HotJava Power | SQL Puzzle | OLAP Tools

DBMS

DATABASE & CLIENT/SERVER SOLUTIONS

VIA AÉREA 02/02/96

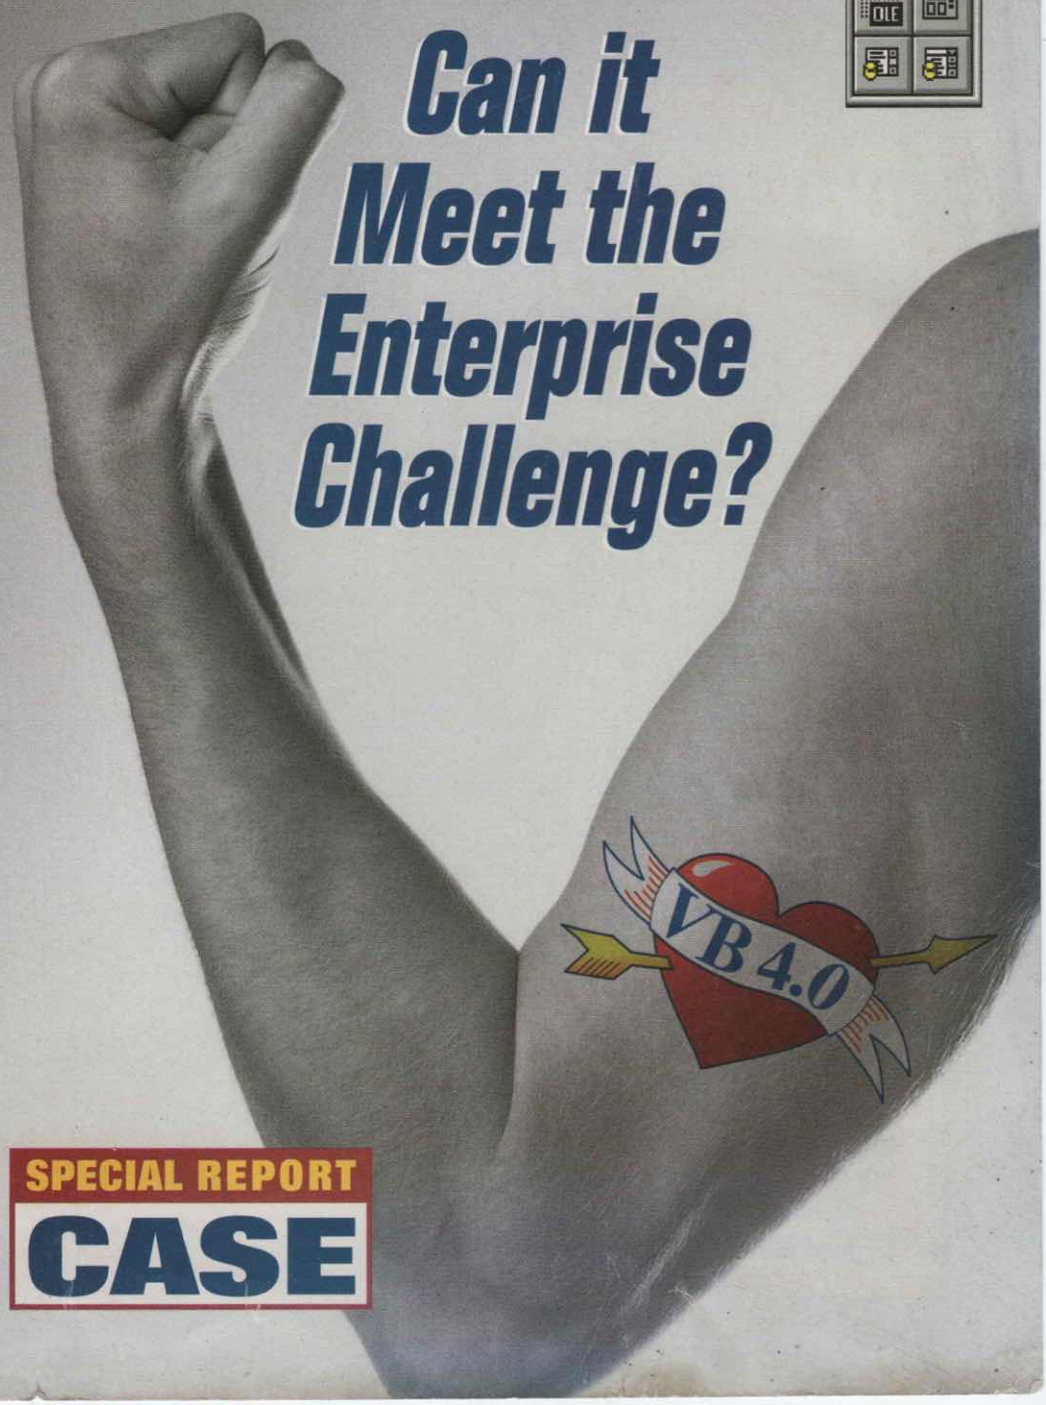
Visual Basic 4.0

Can it Meet the Enterprise Challenge?

- Repository Essentials
- Retrieving Text On Demand
- Predicting Middleware Standards

NEW AND REVIEWED

Magic 6.0, PowerPlay 4.1, Impromptu 3.0



SPECIAL REPORT

CASE

JANUARY 1996 • VOLUME 9 NUMBER 1

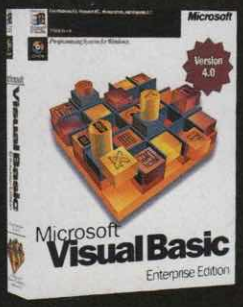
$3.95/$4.95 CAN

01>

0 73361 87392 7

A Miller Freeman Publication

Visual Basic 4.0



*Enterprise-wide multi-tier solutions**

*Partition applications across networks**

*High-performance remote data access**

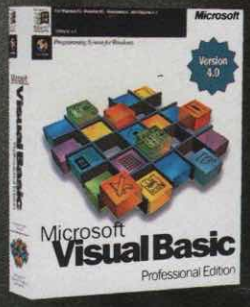
*Large-scale team development**

32-bit with built-in 16-bit support

Open, reusable OLE component creation

Open, integrated development environment

New data-bound OLE controls



Think of it like that guy dave who looks exactly like

your friend dave — talks like him, walks like him — but wait,

this dave **can bench-press a car.**

Introducing Microsoft® Visual Basic® 4.0. Looks like the high-productivity development tool you know, but now enables teams to create multi-tier, scalable, 16- and 32-bit client-server solutions.



Never before has the power of distributed computing been combined with a team-based, rapid application development (RAD) environment for the Windows® operating system. But now, with the Remote Automation technology in the Visual Basic programming system, Enterprise Edition, your development team can quickly build scalable, maintainable applications and deploy them across enterprises big or small — *using a tool they already know.*

Visual Basic 4.0 Enterprise Edition provides fast, direct access to remote client-server databases like SQL Server™ and Oracle.® Plus, its full integration with the Microsoft Visual SourceSafe™ version control system means your development team can securely work in concert — easily sharing and managing code — without ever leaving the Visual Basic environment.

And with full support for Windows 3.1, Windows NT,™ and Windows 95, Visual Basic 4.0 Enterprise and Professional Editions let you transition to 32-bit at your own pace, while preserving your investment in existing code, skills, and technology.** To help shorten development cycles, both versions offer more than 25 new OLE controls, including new Windows 95 controls. Best of all, now you can even create your own open, reusable OLE components such as OLE DLLs and OLE automation servers. So your applications won't just get faster, they'll actually get smarter by design.

To learn more about the Visual Basic 4.0 Enterprise and Professional Editions, call (800) 228-6139, Dept. XU3, for a free Developer Roadmap CD. Or visit **http://www.microsoft.com/vbasic**. For those eager to get started immediately on Visual Basic 4.0 solutions, call one of the resellers below.

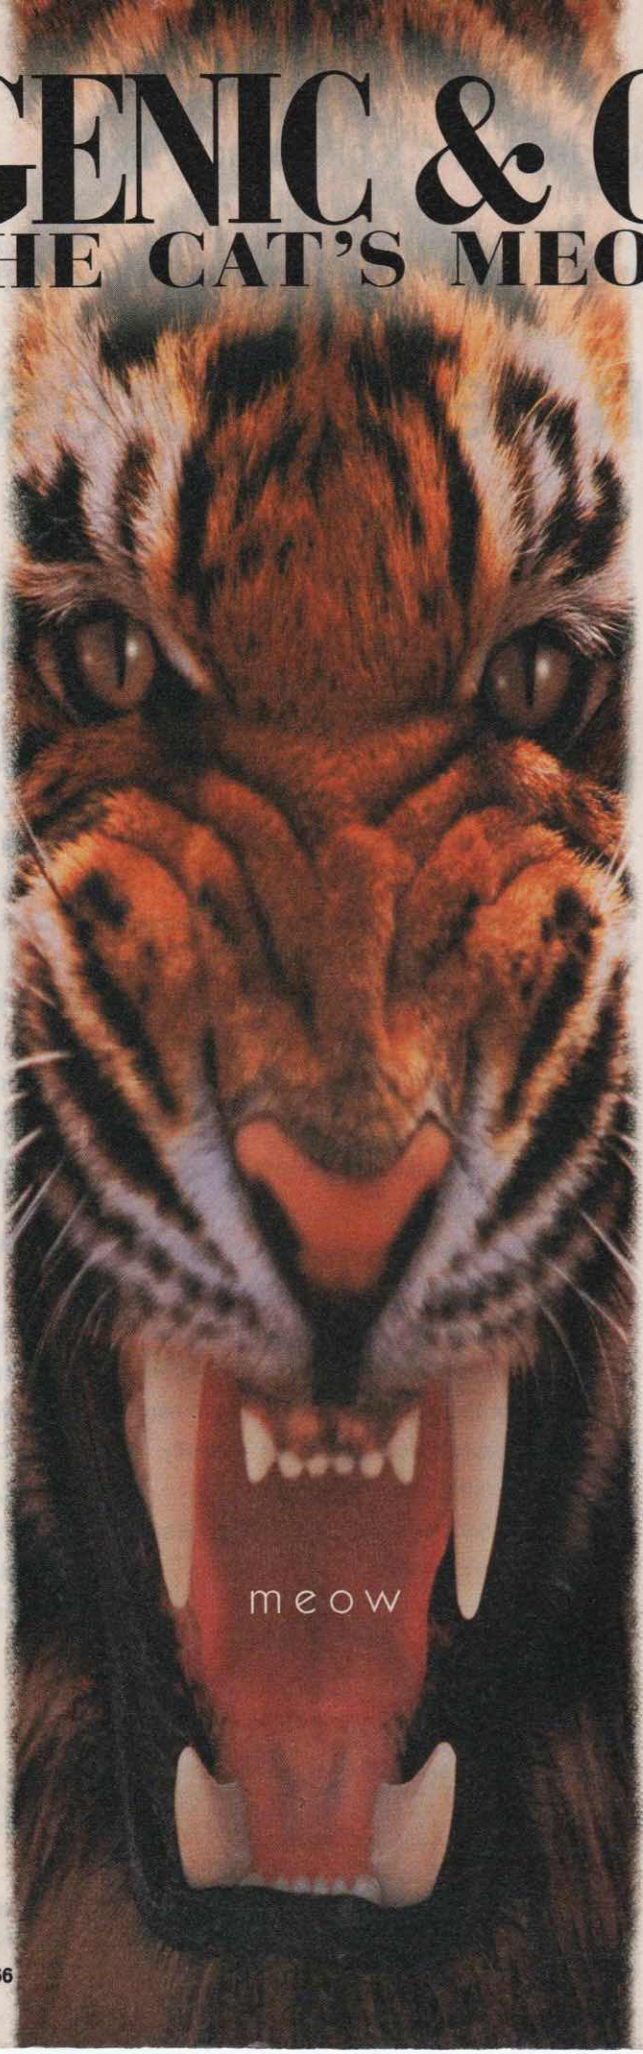
Microsoft

WHERE DO YOU WANT TO GO TODAY?™

*Features available only on Enterprise Edition. **32-bit development requires a 32-bit operating system. Developer Roadmap CD offer expires 12/31/95. ©1995 Microsoft Corporation. All rights reserved. Microsoft, Visual Basic, Windows and Windows NT are registered trademarks and Visual SourceSafe and *Where do you want to go today?* are trademarks of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation.

ASAP Software Express	Programmer's Shop	Programmer's Paradise	Software Spectrum
(800) 883-2911	(800) 446-1185	(800) 892-1187	(800) 840-8194

VISIGENIC & ODBC
THE CAT'S MEOW



Visigenic takes the bite out of your database connectivity problems. From development through deployment, Visigenic's ODBC solutions tear through connectivity barriers. Now you can capture true database independence across Windows, UNIX, OS/2 and Macintosh platforms.

Visigenic has earned its stripes as the number one industry expert in database connectivity. It currently provides ODBC solutions to major industry leaders. Hewlett-Packard. Informix. Microsoft. Oracle. And many more.

ODBC DriverSet

Visigenic ODBC drivers give you the high-speed link between your ODBC-enabled applications and the most popular SQL databases. The Visigenic ODBC DriverSet is composed of drivers, including CA-Ingres, IBM DB2, Informix, Microsoft SQL Server, Oracle, Sybase DBLib, and Sybase System 10.

ODBC Software Development Kits (SDKs)

Visigenic ODBC SDKs let you write and deploy database independent C and C++ applications across any platform. Now you can write applications just once and access any database throughout the enterprise.

Consulting. Training. Support.

Visigenic offers you on-site consulting and training. A broad range of technical support programs designed to fit your needs are also available.

Visigenic is working closely with industry leaders to drive the future of database connectivity technology. To learn more about ODBC solutions from the company with quick reflexes and instant response, contact Visigenic.

Call Visigenic
1-800-632-2864
http://www.visigenic.com
info@visigenic.com

VISIGENIC

YOUR DATABASE CONNECTION

Visigenic and its logo are trademarks of Visigenic Software, Inc. ODBC is a trademark of Microsoft Corporation used under license by Visigenic. All other trademarks are the property of their respective holders.

CIRCLE READER SERVICE NUMBER 66

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

Visual Basic 4.0: Ready for the Enterprise? 44

by David S. Linthicum

With Windows 95 and the subsequent release of Visual Basic 4.0, Microsoft has staked its claim in the enterprise development tool marketplace. This in-depth article reviews VB 4.0, assesses its suitability for the development of strategic enterprise systems, and analyzes VB's new database access, application partitioning, team development, and OLE automation features.

DBMS-Oriented Interactive Voice Response, Part 2 55

by P. L. Olympia, Ph.D.

In the final part of this two-part series, Contributing Editor Pete Olympia explores some alternatives for building computer-telephone integration (CTI) applications, including Periphonics Corp.'s VPS/sp toolset and Voysys Corp.'s VoysAccess System.

Building a DBA Repository System 59

by Martin Rennhackkamp

The purpose of a repository system is to record and manage information about the data models and the designs and implementations of enterprise databases, applications, and systems. But can you apply these application-centric repository principles to physical database management? Yes, you can. This article explains how.

COVER STORY, PAGE 44



January 1996
Vol. 9 No. 1

Cover Art Direction by Barbara Licata

Cover Photography by Sandra Frank

Image Editing by Ron Licata

Thanks to Chas Edwards for lending a hand (and an arm and a bicep).

The Evolution of CASE Tools 67

by Maurice Frank

Most of today's development projects use some kind of modeling tool to design, generate, and manage databases. This article reviews the leading CASE (Computer-Aided Software Engineering) and modeling tools for client/server development, and examines the emerging stratification of the CASE/modeling category into different classes of tools for different kinds of data models: physical, logical, object, and process.

PRODUCTS

Hands-On Reviews...31
Edited by Clara H. Parkes

Magic Software's **Magic 6.0,** *a true "no code" development environment for client/server.*

Cognos Inc.'s **PowerPlay 4.1,** *an analysis tool that emphasizes crosstabs and graphics; and* **Impromptu 3.0,** *a report writer based on templates.*

Soft Notes, *new client/server database products and upgrades.*

COLUMNS

From the Editor
by David M. Kalman 8
Check out *DBMS*'s Buyer's Guide and Client/Server Sourcebook CD-ROM with this issue!

Enterprise C/S
by Judith Hurwitz.....12
The battle rages on for the dominant client/server infrastructure.

Data Warehouse Architect
by Ralph Kimball.....16
Why a freshman in business needs a Ph.D. in SQL to do simple comparisons.

SQL Explorer
by Joe Celko.....20
Joe takes some time out to relate his travel experiences and ANSI work, and provides a new SQL puzzle.

C/S Developer
by David S. Linthicum.....26
Sun Microcomputer's HotJava takes Web-based application development to a new level.

Desktop DBMS
by Tom Spitzer.....84
Text retrieval and management comes of age with next-generation search engines.

DEPARTMENTS

Letters.....6
Event Alerter.....42
Advertisers Index.....72
C/S Connection.....88

COMING NEXT MONTH

The CA Enterprise: An in-depth look at CA's product line, services, and strategies. Plus, how to build a DBA Repository System (part 2), and a comparison of Visual Xbase tools.

DBMS (ISSN:1041-5173) is published monthly, except in June, which is semimonthly and contains the *DBMS* Buyer's Guide Issue, by Miller Freeman Inc., 411 Borel Ave., Suite 100, San Mateo, CA 94402. Second-class postage paid at San Mateo, CA, and additional offices. Entire contents copyright © 1996 Miller Freeman Inc. All rights reserved; reproduction in whole or in part without permission is prohibited. Direct requests for permission to Editor in Chief, *DBMS*, 411 Borel Ave., Suite 100, San Mateo, CA 94402. ***Subscriptions:*** For subscription inquiries, change of address, call: 800-334-8152 (U.S. and Canada), 619-745-2809 (all other countries). Or write *DBMS* at P.O. Box 469039, Escondido, CA 92046-9039. Fax 415-905-2233. Please allow 4-6 weeks for an address change to take effect. Single copy price is $3.95 (Canada $4.95). Subscriptions are $24.95 for one year (13 issues); $43.95 for two years (26 issues) in the U.S. and its territories. Canada (GST# R124771239) and Mexico: $35.00 per year. All other foreign: $60.00 per year. Foreign orders must be prepaid in U.S. Dollars on a U.S. bank. Canada Post International Publications Mail Product (Canadian Distribution) Sales agreement No. 0548669 ***Postmaster:*** Send address changes to: *DBMS*, P.O. Box 469039, Escondido, CA 92046-9039. ***Foreign Newsstand Distributor:*** Worldwide Media Service Inc., 30 Montgomery St., Jersey City, NJ 07302; 201-332-7100.



"So Risk-Free, It's A No-Brainer."



The INTERSOLV Virtual Data Warehouse begins at the user's desktop, providing immediate access to enterprise data. No complex infrastructure or costly set-up is needed. This dynamic, demand-driven approach is open, fast, reliable and leverages any existing environment.

The Virtual Data Warehouse is implemented through the plug-and-play components in the INTERSOLV DataDirect Series. With them, you'll reap the benefits of a data warehouse immediately. And users will get the data they need instantly.



From $199

DataDirect SmartData is built on our leading ODBC driver technology and shields users from the complexities of SQL, joins, and cryptic column names. It provides quick access to data and delivers it in business terms. And because SmartData is an open architecture, you can use it with your favorite tools and applications.

DataDirect Explorer is the presentation component of the Virtual Data Warehouse. It lets users quickly and easily query and analyze data and create charts, reports, graphs and forms without MIS support.



From $499

And when you combine Explorer with SmartData, the complexities of databases disappear.

Less time. Less hassle. Faster results. Choose the rapid approach to data warehousing. Give users a way to get data when, where and how they want it in everyday terms. Order DataDirect SmartData and Explorer today. A free demo CD and white paper are available upon request.



INTERSOLV

800-876-3101 ext. S808

http://www.intersolv.com

CIRCLE READER SERVICE NUMBER 86

LETTERS • LETTERS • LETTERS

Farmer Faux Pas

Obviously Mr. Labbauf, one of those who contributed an opinion about Windows 95 in your November issue (From the Editor, page 89), is a stranger to farms. His expression "bail and wire and bubble gum" should be "baling wire and bubble gum," a well-known rural engineering technique.

Sincerely,
Patrick J. Brennan
Virginia Beach, Va.

The Debate Continues...

I refer to the October 1995 letter from Joe Celko. The principal technical disagreement between Celko and myself (at least, in the area under discussion) concerns the difference between *types* and *representations*. Let me lay a little groundwork:

- A *type* — usually called a *domain* in relational contexts — is a named set of *values*. Every type has an associated set of *operators* that apply to values of that type.
- Given a type, a *representation* — also known as a *structure* — for that type is an *encoding in memory* for values of that type.

Thus, types are a *model* consideration, while representations are an *implementation* consideration. As a trivial example, consider the type FLOAT in SQL. The representation of FLOAT numbers is not specified in the SQL language, and users don't need to know what it is; what they do need to know is the operations they can perform on FLOAT numbers — "+", "*", ">", and so on. And that's *all* they need to know. To quote Luca Cardelli and Peter Wegner ("On Understanding Types, Data Abstraction, and Polymorphism," *ACM Comp. Surv. 17*, No. 4, December 1985): "A major purpose of type systems is to avoid embarrassing questions about representations, and to forbid situations in which these questions might come up."

Now I turn to Celko's letter. At one point, Celko seems to agree with the foregoing, when he says that "numbers...and numerals...are [not the same thing]" — that is, numbers are the type, numerals are the representation. Very confusingly, though, he uses the term *type* to mean *representation* — which is presumably why he also says "I do not see how domain and data type can [possibly] be the same." The source of this unorthodox terminology is probably the SQL language, which is thoroughly confused in this whole area (my FLOAT example notwithstanding).

However, Celko also gives an extensive example that suggests that he does *not* properly understand the type *vs.* representation distinction after all: "I want to represent [sic] temperatures. Degrees Kelvin is one possible scale. I could use Celsius or Fahrenheit instead; they all use numeric values. Or I could use the scale ('hot,' warm,' 'cool,' 'cold'); the data type is character string. What is the right choice? [Given a certain application area,] I will probably prefer Kelvin because it is expressed with a numeric data type. I can use arithmetic operators with numeric data types, and I will need to do calculations. [But] you cannot add temperatures..." (somewhat edited and condensed from Celko's letter).

I have two major observations on this example. The first has to do with the type *vs.* representation distinction and is fundamental. The second is not fundamental but has to do with a question of good design.

1. The data type is *not* character string for "the scale ('hot,' 'warm,' 'cool,' 'cold')"; the *representation* is character string. We're not going to perform substring or concatenate operations on temperatures! Likewise, it's the *representation*, not the data type, that is numeric for the Kelvin and Celsius and Fahrenheit designs; as Celko himself points out, we're not going to add temperatures.

   No: The operations we're going to perform on temperatures are, precisely, the operations that are defined for the temperature *type*, whatever those might be.
2. Celko also asks which of his possible designs is the right choice. Presumably he's saying we could have a domain called, say, K_TEMP, whose values are measurements on the Kelvin scale; or a domain called, say, C_TEMP, whose values are measurements on the Celsius scale; or several other analogous possibilities.

Personally, I think *none* of these proposals is the best choice. Rather, I would have a single "temperature" domain, with operators to expose temperature values as "number of degrees Kelvin," "number of degrees Celsius," "number of degrees Fahrenheit" — even as "hot," "warm," and so on, if desired. One advantage of this design is that it avoids the need for explicit conversion functions from one scale to another (any such conversions would be completely under the covers). Another is that it affords a greater degree of data independence (we can change the representation from, say, Kelvin to Fahrenheit with logical impunity).

Celko makes many further points in his letter with which I strongly disagree. If I went into detail on all of them, however, this would be a very long letter. Let me therefore just make a few apparently dogmatic assertions here, without trying to provide any evidence in support of them (I can provide such evidence on request):

- Celko's remarks on "viewing each entity-attribute combination as a separate domain/data type" are nonsense.

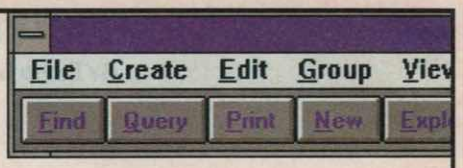
(continued on page 76)

WEB INFO AND OTHER IMPORTANT STUFF

☞ Check out *DBMS* on the World Wide Web for instant, 24-hour access to a rich blend of news, product information, selected articles, and resources for database professionals. **Just point your Internet browser to http://www.dbmsmag.com.** Our site includes: the *DBMS* Buyer's Guide & Client/Server Sourcebook; a calendar of database-related conferences, trade shows, and seminars; subscriber service and frequently asked questions (FAQs); late-breaking client/server news; *DBMS* interviews and editorials; writer's guidelines; and online letters to the editor.

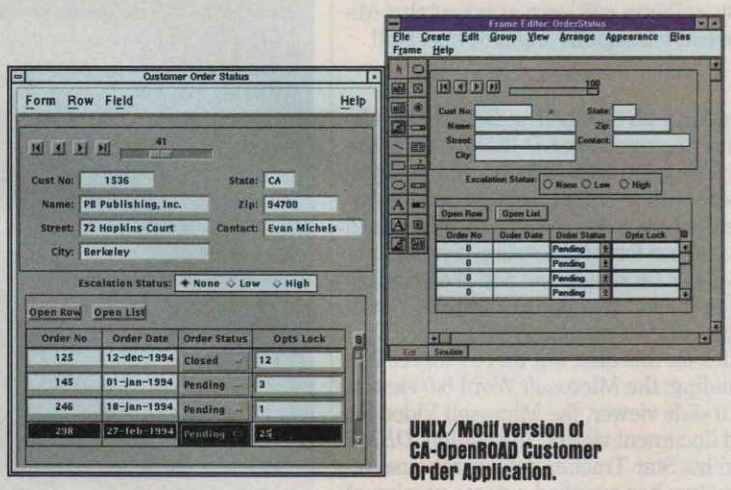
☞ *DBMS* welcomes letters to the editor. We encourage you to respond to any editorial content, and to discuss and critique industry issues and trends. We also welcome suggestions of what you would like to see in future issues of *DBMS*. Letters must include your name, address, and telephone number.

☞ Mailing address: *DBMS* Letters, 411 Borel Ave., Ste. 100, San Mateo, CA 94402. You can also send letters via fax at 415-358-9855, Internet mail (dbms@mfi.com), or CompuServe (in the *DBMS* Forum or at 76711,466). *DBMS* reserves the right to edit letters.



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





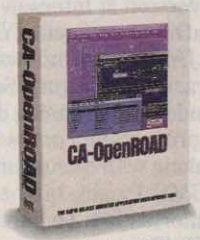
UNIX/Motif version of CA-OpenROAD Customer Order Application.

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. 13301.

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.

CIRCLE READER SERVICE NUMBER 188

BY DAVID M. KALMAN

A Slipped Disc



Perhaps you saw it Sunday, November 5 on CBS. A NASDAQ ad equated Bill Gates with Thomas Edison, claiming that Gates "invented" the software that runs our computers. Sorry NASDAQ, but not even close. Gates may be a visionary and a marketing genius, but he's no Thomas Edison. Operating systems and computer applications existed for years before Gates was even born. Gates applied existing technology to new markets, and to a new platform: the microcomputer. Edison, in contrast, invented several new "applications" of basic electro-mechanical principles, including the incandescent light bulb (1879) and the phonograph (1877). Next thing you know, children will learn in school that Microsoft invented the GUI and the use of the mouse as a computer input device. Not!

Install and Run Our New CD-ROM

Speaking of inventions, you'll find a nifty one slipped in with this issue of *DBMS*. I'm talking about our new Buyer's Guide and Client/Server Sourcebook CD-ROM. This is the second edition of this popular disc, and this year it's loaded with software and information. The disc contains our popular Buyer's Guide, both in HTML and executable formats, and a database of more than 43,000 U.S. five-digit zip codes. Also on the disc, you'll find a copy of the entire *DBMS* Web site that you can use for local viewing. If you do not have a Web browser, we provide Netcom's NetCruiser (in standalone and Web-capable versions). For accessing CompuServe, we provide the latest version of Dvorak Development's NavCIS Pro, preconfigured to access the *DBMS* CompuServe Forum.

I'm getting a bit winded listing all of the goodies on this disc, but there's more. The disc includes a variety of document viewers, including: the Microsoft Word 6.0 viewer, the Adobe Acrobat 2.1 viewer, the PowerPoint 4.0 slide viewer, the Microsoft Video for Windows 1.1 Runtime, and the Common Ground document viewer. In addition, *DBMS* Contributing Editor Ralph Kimball has provided his Star Tracker data warehouse reporting tool, and Great Lakes Business Solutions Inc. has supplied a demo version of its Wise Installation System (which we used to create the disc's setup program).

The CD-ROM also contains software and info from advertisers who would like very much to get your attention. Btrieve Technologies provides spec sheets and white papers on the company's products; Cast Software provides a demo of its repository-based development tools; IBM offers a self-running demo of DB2 version 2; Intersolv provides product white papers and a self-running demo of its DataDirect products; JYACC includes product info in HTML format and a demo version of JAM 7 for Windows; Oracle provides a complete Personal Oracle7 desktop DBMS, Oracle7 Workgroup Server for Networks, and Oracle Power Objects; Popkin Software includes a self-running demo of the System Architect CASE tool; Sybase offers an HTML presentation on its Portfolio, PowerBuilder, S-Designor, and Watcom SQL products, along with evaluation versions of S-Designor Professional 4.2 and the Watcom SQL DBMS. Finally, on the disc you'll find some *DBMS* writer's guidelines and subscription information. To purchase additional copies of the CD-ROM, contact Gae Reck at 415-655-4245.

DBMS Tech Editor Maurice Frank handled the production of the CD-ROM, coordinating the software vendors, developing the setup routine, and writing the readme files, all in his spare time. Of course, Maurice had the help of his cats, Spock and Micro, whose images you'll find in the /dbmswww directory. Way to go Maurice, Spock, and Micro.

Surfing for Dollars

I almost forgot to mention that 10 copies of the CD-ROM include a secret button hidden in one of the advertisers' demonstrations or documents. The first eligible person to activate the secret key and return his or her disc will receive $500 (and a replacement CD-ROM). The next nine eligible winners to send in their discs will receive $50 each. ■

DBMS

Publisher
Philip Chapnick

EDITORIAL

Editor in Chief David M. Kalman
Managing Editor Kathleen B. O'Connor
Features Editor Theresa Rigney
Associate Editor Clara H. Parkes
Technical Editor Maurice Frank
Contributing Editors Shaku Atre, Joe Celko, Judith Hurwitz, Ralph Kimball, David Linthicum, Stewart McKie, P. L. Olympia, Tom Spitzer, Jeff Winchell

ART AND PRODUCTION

Art and Production Supervisor Barbara Licata
Technical Illustrator Margaret Anderson

ADVERTISING

Eastern United States Len Ganz
Account Representative Paul Calento
Western United States Susan Dawson
Account Representative Joseph DeSantis
Ad Coordinator Gae Reck
Advertising/Editorial Assistant Connie Stevens

MARKETING

Marketing Account Manager Sue Singh
Marketing Production Designer Lynn LaRocca

CIRCULATION

VP/Director of Circulation Jerry Okabe
Group Circulation Manager Kathy Henry
Circulation Manager Lisa Eversole
Assistant Circulation Manager Phil Payton
Newsstand Manager Pam Santoro

Vice President/Group Director
Steve Schneiderman

Miller Freeman
A United News & Media publication

Chairman/CEO Marshall W. Freeman
Executive Vice President/COO Thomas L. Kemp
Senior Vice President/CFO
Warren "Andy" Ambrose
Senior Vice President David Nussbaum
Senior Vice President H. Verne Packer
Senior Vice President Donald A. Pazour
Senior Vice President Wini D. Ragus
Vice President/Production Andrew A. Mickus
Vice President/Circulation Jerry Okabe

PRINTED IN THE USA

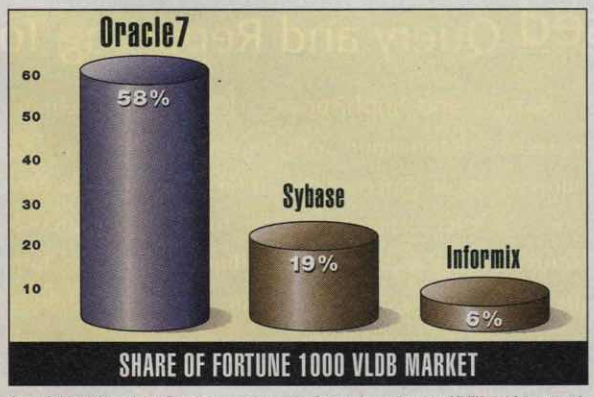
BPA INTERNATIONAL

For subscription orders or change of address: call 619-745-2809 or toll free 800-334-8152, write *DBMS*, P.O. Box 469039, Escondido, CA 92046, or contact our customer service department via CompuServe at 71572,341.

Data Warehouse Comparison

Data warehouses are central information repositories for storing and analyzing vast amounts of historical and reference data from a number of different sources. Oracle7 has many more features for data warehousing than any other DBMS.

Features	ORACLE7	Sybase System 10	Informix 7.1
Production Terabyte Databases Supported	Lots	Zero	None
Scalable SMP Support	Yes	Nope	Yes
Scalable MPP Support	Yes	No Way	No
Parallel Query	Yes	Nope	Yes
Parallel Load	Yes	Not	No
Parallel Index	Yes	No	Yes
Fault-Tolerant Parallel Server	Yes	Never	No
Parallel User-Defined Functions	Yes	Nope	No
Bit-Map Indexes	Yes	Yes	No
Star Query Optimization	Yes	No Way	No
Integrated Replication	Yes	Yes	No
Multidimensional Cache	Yes	Sorry	No
Video Storage & Playback	Yes	Nyet	No
Text Storage & Search	Yes	Not	Yes
Spatial Data	Yes	Never	No



Study fielded 6/26–9/1/95. Based on a census of the F1000 who are running UNIX and have a single DB ≥ 50GB. Source: Griggs Anderson Research.

The Oracle data warehouse provides fast query access to all your data: relational, spatial, text, images, audio, video and instant multidimensional drill-down analysis. So, if fast, accurate decisions are important to your company's performance, call Oracle at 1-800-633-0509, ext. 6897. Or download videos explaining the Oracle Warehouse at http://www.oracle.com/info/dw/video.html.

ORACLE®

Enabling the Information Age

©1995 Oracle Corporation. All rights reserved. Oracle is a registered trademark. Oracle7 is a trademark of Oracle Corporation. All other product names are property of their respective owners.



Introducing IQ/Objects.™

The Only True Object-Based Query and Reporting Tool.

You already know that object technology is revolutionizing application development. Now you can put the power of objects in the hands of every business user.

A New Paradigm. With IQ/Objects, users create queries and reports with prefabricated building blocks called SmartObjects.™ SmartObjects are intelligent. They communicate with each other. And they're reusable. Users simply assemble these SmartObjects to create a limitless array of sophisticated queries and reports.

Simple and Sophisticated. IQ/Objects integrates and correlates information with unprecedented ease. With no more effort or expertise than it takes to use a word processor, any user can...query data from heterogeneous data sources...combine charts, multidimensional cross-tabs, and free form layouts on a single page...use parent/child reports to pass results from one query object to another...make any query object conditional...add watermarks, bitmaps, even audio and video clips!

IQ Software Corporation (770) 446-8880, fax: (770) 446-0321, email: sales@iqsc.com. World Wide Web: http//www.iqsc.com. European Subsidiaries & Affiliates: UK (01962) 844777, France (01) 474-90304, Italy-Dexia Italia (039) 6-517933. International Distributors: Australia (02)369-1932, Mexico 5-575-6098 (D.F.) 83-490-172 (Mty.), Singapore 227-7988, South Africa (011) 421-4800, Spain 50-27-55-11, Taiwan 2-740-8020.

See Query and Reporting In A Brilliant New Light.

"A breakthrough technology for query and reporting tools. A must on every enterprise's short list for data access solutions."

Bob Moran
Aberdeen Group

"IQ Software has brought the power of objects to mainstream business intelligence."

John Rymer
Patricia Seybold Group

"We looked at them all. IQ/Objects is indeed the best."

Mitch Kaufman
Finova

Usability and Scalability. A Microsoft Office look and feel makes IQ/Objects instantly usable. And, by combining IQ/Objects with IQ SmartServer for UNIX or NT, you get a highly scalable, three-tier architecture solution – ready for everything from workgroup decision support to enterprise-wide data warehouses.

Contact us today for your FREE evaluation copy of IQ/Objects. Compare its breakthrough technology against any other query and reporting tool. The difference is night and day.

IQ SOFTWARE®

800-458-0386

sales@iqsc.com

http://www.iqsc.com

For distributors in other countries, call IQ Software's U.S. headquarters. © 1995 IQ Software Corporation. The IQ Software logo is a registered trademark and IQ, IQ/Objects and IQ SmartServer are trademarks of IQ Software Corporation. All rights reserved. All other product names are trademarks or registered trademarks of their respective companies.

CIRCLE READER SERVICE NUMBER 116

BY JUDITH HURWITZ

Battle for the C/S Infrastructure



WINNING THE BATTLE (AND THE WAR) FOR A STANDARD MIDDLEWARE INFRASTRUCTURE FOR CLIENT/SERVER COMPUTING.

A battle is brewing for dominance of the standard middleware infrastructure for client/server computing. Ever since it became clear that the mainframe environment would not continue to be dominant, there has been confusion about what would replace it. During the heyday of the "open system" movement, it appeared that a consensus might form around products such as Open Software Foundation's (OSF's) Distributed Computing Environment (DCE). It is increasingly clear to me, however, that the de facto infrastructure will emerge out of a fierce battle for control. Who will win this battle? There is no simple answer.

In this two-part series, I suggest four possible scenarios for the future of the client/server infrastructure, and offer my assessment of each scenarios' probability of success. I discuss the first two this month and the final two next month. Each of these four scenarios will likely emerge over the next five years, and each has the potential to succeed. I'd also like to hear your opinions. Feel free to email me via the Internet at Jhurwitz@world.std.com.

Four Scenarios: Pick Your Poison

The Super-Packaged Software Scenario. Suppose a company in the packaged software market begins to dominate the market for several key vertical segments. Developers for this company have written their own comprehensive middleware products to support the corporate applications. Other developers begin to use these proprietary tools to create other specialized applications. Over time, this infrastructure becomes a default standard for a large segment of the Fortune 1000 and then trickles down to become a standard for a broad range of corporations.

The Microsoft Scenario. Microsoft (Redmond, Wash.) has amassed enough clout and mindshare to set the agenda for a future client/server infrastructure. The company successfully stalled the evolution of a standard middleware infrastructure for several years by promising a client/server version of Object Linking and Embedding (OLE). This gave Microsoft the time to rewrite OLE so that it would be appropriate as a large-scale client/server infrastructure. Microsoft uses the resources of the most successful, large software and integration companies, such as Digital Equipment Corp. and Software AG, to help it add the additional pieces it needs to complete an enterprise infrastructure. Microsoft is now poised to be the IBM of the early 21st century.

.....

Judith Hurwitz is president of Hurwitz Consulting Group Inc. (Newton, Mass.), a consulting, publishing, and research services firm specializing in client/server development tools, client/server infrastructure, and systems management. You can reach Judith at 617-965-6900 or via the Internet at Jhurwitz@world.std.com.

The United Front Scenario. Several large system vendors and partners, including systems integrators and large, independent software companies, come together and agree on certain infrastructure components that they all promise to use. These vendors are motivated by fear of domination by either the packaged software players or Microsoft. ISVs flock to this solution because it protects them from total dependence on either of these powerful forces.

The Internet-as-Infrastructure Scenario. The Internet, with its flat hierarchy and directory structure, forms the foundation for a new infrastructure. ISVs, sensing opportunity, add a variety of low-level services, such as tight security and development tools, to turn the Internet into a true client/server infrastructure. This is very appealing to customers because no one company prevails.

And the Winner Is?

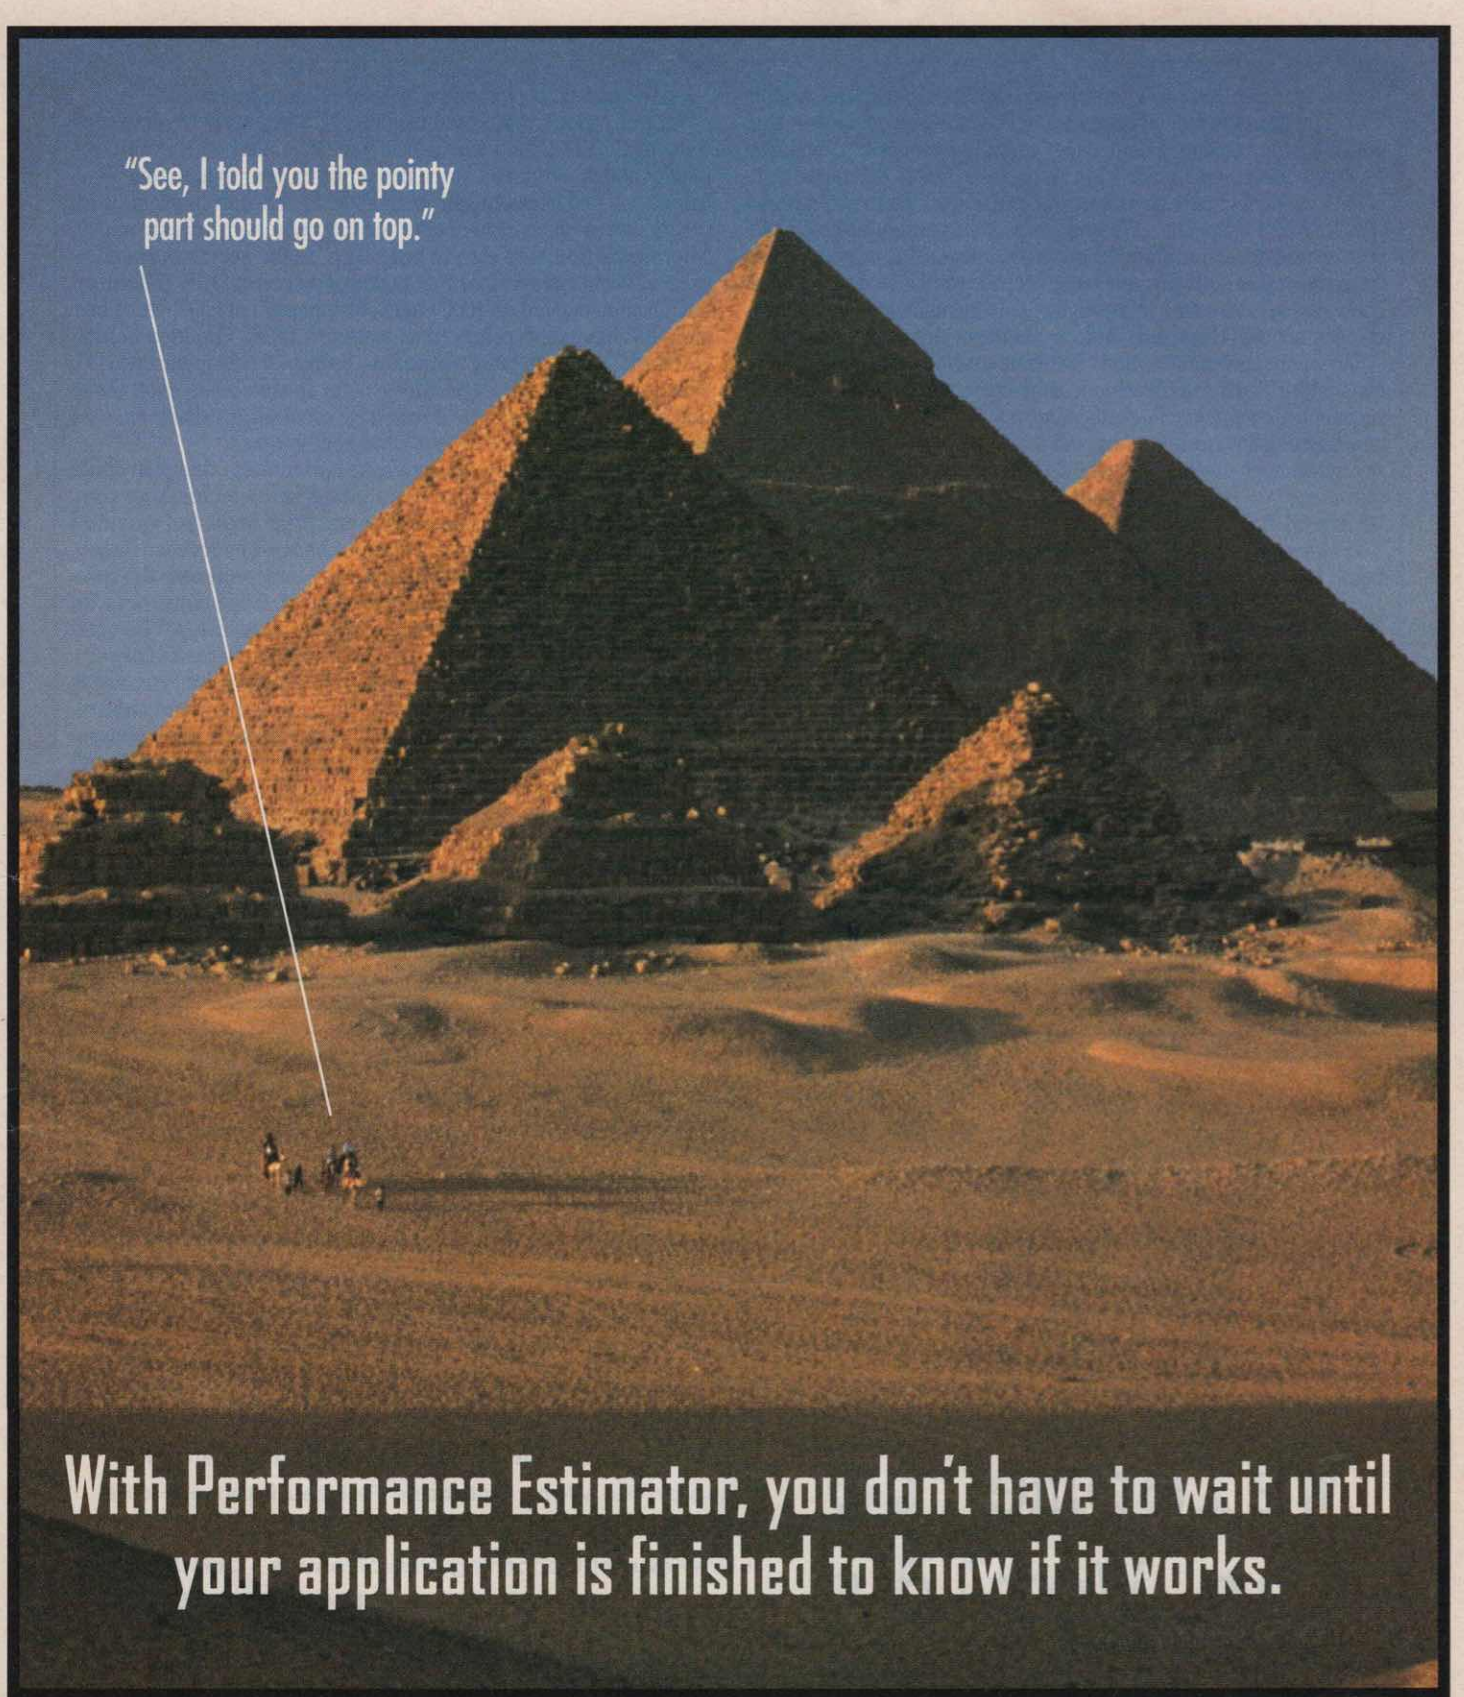
Which of these scenarios is most likely to succeed? The answer is far from obvious. In fact, each scenario has a chance of winning, depending on several other factors. For example, will the Internet ever be truly secure? Will Microsoft be able to master the server side of the computing business? Keep these questions in mind as I assess each scenario and its chance of success, based on certain assumptions. If an assumption turns out to be wrong, all bets are off.

The Super-Packaged Software Scenario

Probability of success: 60 percent

In order for this scenario to succeed, the following things would have to happen: A packaged software vendor would supply a series of well-designed vertical and horizontal software packages that would meet a large percentage of corporate requirements. This means that the vendor would provide a variety of financial software modules, human resources software, and industry-based vertical software (manufacturing, health care, and so on). I am assuming that this software would be extremely modular and flexible. It would therefore be relatively easy for a company to continue to add modules and bring in software from outside sources to fill gaps. It would also be exceedingly easy to provide integration among these functions.

In addition, I am assuming that the company would use de facto and de jure standards to create both a remote procedure call (RPC)-based and messaging-based infrastructure, because support for both messaging and RPC would be necessary to meet a broad range of development requirements. This company would also have to create a comprehensive set of systems and applications management tools designed to help organizations master the deployment side of client/server computing.



With Performance Estimator, you don't have to wait until your application is finished to know if it works.

As the ancient Egyptians used to say, "The only way to ensure your application's performance is to model it first." They may have been on to something. And new PLATINUM Performance Estimator makes modeling so easy you can test out even your most imposing projects. Performance Estimator is extremely rich in functionality, providing over 25 reports that give you critical insight into run/response time, physical and logical I/Os, CPU workloads and system contention. So you won't have to spend valuable time and energy building an application that may not stand the test of time. Email us at info@platinum.com for a comprehensive WHITE PAPER or call 1.800.442.6861. PLATINUM PERFORMANCE ESTIMATOR

PLATINUM
TECHNOLOGY
The Open Enterprise Management Company

PLATINUM *technology, inc.*, 1815 South Meyers Road, Oakbrook Terrace, Illinois 60181 Phone 708.620.5000 Fax 708.691.0718 Web http://www.platinum.com All product names mentioned are trademarks or registered trademarks of the respective owners and their subsidiaries. ©1995 PLATINUM *technology, inc.*

CIRCLE READER SERVICE NUMBER 57

Another assumption I would make for this scenario is important: The company must understand how to create a toolset that enables other software companies to build new applications that leverage the environment's infrastructure. If this packaged software vendor is good at partnering with other software companies, it will publish its APIs and make it extremely easy for others to add functionality. If enough influential corporations implement this software environment, they will also demand that their key vendors and partners use this infrastructure. Systems integrators will then sign on to provide service and support.

What are the ramifications if this scenario is successful? The most significant ramification is that there would be a profound impact on operating system vendors. If you assume that the packaged software company will provide a cross-platform environment, then operating systems, databases, and even systems platforms will take a back seat. This company would therefore be in a position to dictate standards in several different areas. SAP America Inc. (Philadelphia), with its Basis middleware layer, is one company that could conceivably pull this off.

The Microsoft Scenario

Probability of success: 40 percent

Microsoft continues to urge customers to delay implementing emerging infrastructure components such as messaging middleware, Common Object Request Broker Architecture (CORBA), DCE, and other technologies because it promises to provide a straightforward, clear interface from desktop to enterprise server based on OLE. However, its initial server implementations of OLE are very limited in functionality. OLE objects can be physically stored on servers, but little else is accomplished. However, many users are excited about a Microsoft-based infrastructure and are willing to wait. They believe that Microsoft's market clout makes it a likely victor in the war to dominate the computing infrastructure. To some users, this approach seems more practical than trying to piece together an infrastructure that might fail.

Organizations that spent millions piloting early implementations of CORBA and DCE are now very familiar with their complexity and would love to turn the problem over to someone else. Microsoft has now bought itself several years to begin putting the pieces in place to take over the client/server infrastructure. Microsoft's first step is to rewrite OLE on the server into a comprehensive set of services that includes pieces of DCE (Microsoft has already copied its RPC) and components of CORBA to create what Microsoft refers to as Network OLE. The idea behind Network OLE is to create seamless links to desktop-based OLE. The company brings in partners eager to ride on its coattails to design the complex tools needed by most enterprises, such as security, performance monitoring, and so on.

In theory (and given enough time and money) this plan could work. Microsoft will probably purchase some technology companies that have expertise at the enterprise level. So, why do I give this scenario only a 40 percent chance of success? Essentially, I think that Microsoft will seriously underestimate the complexity and demands of getting the job done. It will come out with interim deliverables that will fall short of its promises. The complexity of distributed computing will most likely make corporations realize that a large-scale, general-purpose infrastructure is far too complex for any one company to implement and control.

What are the ramifications for the industry? If this scenario fails to materialize, the computer industry will relegate Microsoft to a lesser role than the software giant would like. However, Microsoft will clearly continue to dominate the desktop platform. I expect that the company will do exceedingly well in the LAN-based database market. In addition, Microsoft will continue to sell more and more component products that leverage its desktop franchise. This will open the door for more healthy competition and innovation. Next month, I'll discuss the final two scenarios: The United Front and the Internet-as-Infrastructure. ■



Welcome to the World's Fastest DBMS!

We've put *DBMS* on the Web, a free electronic resource that supplements our monthly magazine. Just point your internet browser to http://www.dbmsmag.com and you'll get instant access to a rich blend of news, product information, selected articles, and resources, including:

• *DBMS* Buyer's Guide & Client/Server Sourcebook • Calendar of Database-Related Conferences, Trade Shows, and Seminars • *DBMS* Subscriber Service and Frequently Asked Questions (FAQ) • Late-Breaking Client/Server News • Database Vendor Sites on the Web • Database-Related Internet News Groups • *DBMS* Interviews and Editorials • Previews of Upcoming *DBMS* Issues • *DBMS* Writer's Guidelines and Online Letters to the Editor

http://www.dbmsmag.com

Wait No More — *DBMS* on the Web

Want to see more? Don't wait another minute. Load up your Web browser and join the world's fastest DBMS – *DBMS* on the Web!

Power Objects.
Because every line of code you write ends up on your face.



Unique drag and drop technology dramatically simplifies applications development.

$99
Introductory Offer
Personal Edition
Over $150 Savings*



Unnecessary and repetitive code takes its toll. On your business, on your mind, and, yes, even your face. But relax. Apple®, the experts in easy to use technology, has teamed up with Oracle, the database experts, to bring you a winning business solution. Oracle Power Objects™.

A visual programming tool, Power Objects automates complex database functions by letting you reuse code quickly via drag-and-drop object technology. With Power Objects, you can insert and delete data located on a server DBMS with full transaction support and no coding whatsoever. The same task in Visual Basic would require many redundant lines of code. You can even run Windows app's unmodified on a Macintosh™ and vice versa. And Power Objects uses a scripting language you already know, standard BASIC.

To eliminate unnecessary lines of code from your programming, and perhaps elsewhere, buy Oracle Power Objects today, before it's too late. Available for both Macintosh and Windows in the Personal or Client/Server Edition.

	Power Objects	Visual Basic
Drag-and-drop database programming	Yes	No
Integrated SQL database	Yes	Yes
Industry standard BASIC language	Yes	Yes
Seamless scalability to Oracle 7 & SQL Server	Yes	No
Native Oracle 7 & SQL Server interfaces	Yes*	No
True object-oriented code reusability	Yes	No
OCX and OLE 2 support	Yes	Yes
Multiplatform: Windows, Macintosh, OS/2**	Yes	No

** Standard in the client/server version*
*** Scheduled for release in early '96*

Limited time offer. Order today by calling Apple direct at
1-800-950-5382 ext. 906,
visit the Oracle Store™ at
http://www.oracle.com
or see the Apple reseller nearest you.

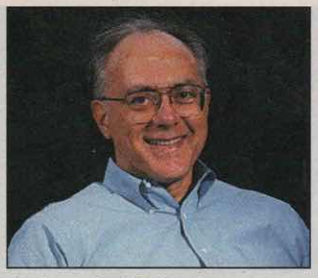
Mac OS For Windows

ORACLE
Enabling the Information Age

Offer valid through March 31, 1996. Tax, shipping and handling costs are not included. *©Apple Computer, Inc. All rights reserved. Apple, the Apple logo, the Mac OS logo, and Macintosh are registered trademarks of Apple Computer, Inc. in the U.S. and other countries. Oracle Power Objects and Oracle Store are registered trademarks of Oracle, Corp. All other trademarks are the property of their respective owners.* *Prices listed are special promotional prices and reseller prices may vary. Offer valid only in the U.S. and not valid in connection with any other offers. *Based on the regular estimated retail price.*

By Ralph Kimball

The Problem with Comparisons



A FRESHMAN IN BUSINESS NEEDS A PH.D. IN SQL, PART 1.

The most difficult area of data warehousing is the translation of simple business analyses into SQL. SQL was not designed with business reports in mind. SQL was really an interim language designed to allow relational table semantics to be expressed in a convenient and accessible form, and to enable researchers and early developers to proceed with building the first relational systems in the mid-1970s.

How else can you explain the fact that there is no direct way in SQL to compare this year to last year? Or to compare one group of customers to another? Or to number the answer set? Very simple business questions require complex and daunting SQL. In this and next month's columns, I explore the two most glaring SQL issues: the inability to perform comparisons and the lack of sequential computations. I attack the comparison issue this month.

The early days of relational systems were marked by arguments at conferences between the academics at Berkeley and members of the more traditional Cobol camp. When the Cobol camp withdrew, some of the standard techniques that had been learned from data processing were lost in the shuffle. This situation would not have been too bad if SQL had been more dynamic and alive as a language over the intervening 20 years, but the efforts of the language committees for both SQL-89 and SQL-92 were concerned with extending the relational semantics of SQL, rather than with adding simple business data analysis capabilities. Therefore, the industry is struggling today with a 1970s interim relational processing language that was invented by kindly academic professors, rather than by businesspeople.

Let me present some examples from common business reports to show what is so hard in SQL. The first, and most glaring, problem is the difficulty of performing simple comparisons. Suppose you want to show this year's sales versus last year's sales:

	4Q95 Sales	4Q94 Sales
Doodads	57	66
Toggles	29	24
Widgets	115	89

.....

Ralph Kimball was coinventor of the Xerox Star workstation, the first commercial product to use mice, icons, and windows. He was vice president of applications at Metaphor Computer Systems, and founder and CEO of Red Brick Systems. He currently designs large data warehouses for a living. His forthcoming book, The Data Warehouse Toolkit: How to Design Dimensional Data Warehouses *(Wiley), covers data warehouse design techniques. You can reach Ralph on the Internet via his home page at http://www.rkimball.com.*

This little report is the bedrock of business analysis. A number in the business world rarely makes much sense unless it is compared to some other number. In this report, the comparison numbers are side by side, but you could just as well have used them together in a ratio to show the growth in sales from last year to this year. This-year-versus-last-year is the standard example. You can see that sales of Widgets and Toggles are up substantially from last year, and the sales of Doodads are down. This immediately suggests further analysis and perhaps even management action.

Amazingly, it is very difficult in SQL to get an answer set that looks like this. The usual SQL template fails you. You start by writing a fragment of the SQL needed for the fourth quarter 1995 sales (4Q95):

```
SELECT Product.Product_Name, SUM(Sales.Dollars)
FROM Sales, Product, Time
WHERE . . . Time.Quarter= '4Q95'
AND Product.Product_Name in ('Doodads', 'Toggles',
 'Widgets')
GROUP BY Product.Product_Name
```

If you have been reading this column for the past few months, you would expect that the underlying dimensional database schema consists of a large central fact table (Sales) joined to a few dimension tables (including Product and Time). To keep this column focused, I have omitted some of the bookkeeping required in the SQL, such as listing all the join constraints in the WHERE clause required to hook the tables together.

Now, how do you get the sales numbers for 4Q94? SQL leaves you with no good alternatives. If you expand the Time constraint to include both quarters, as follows:

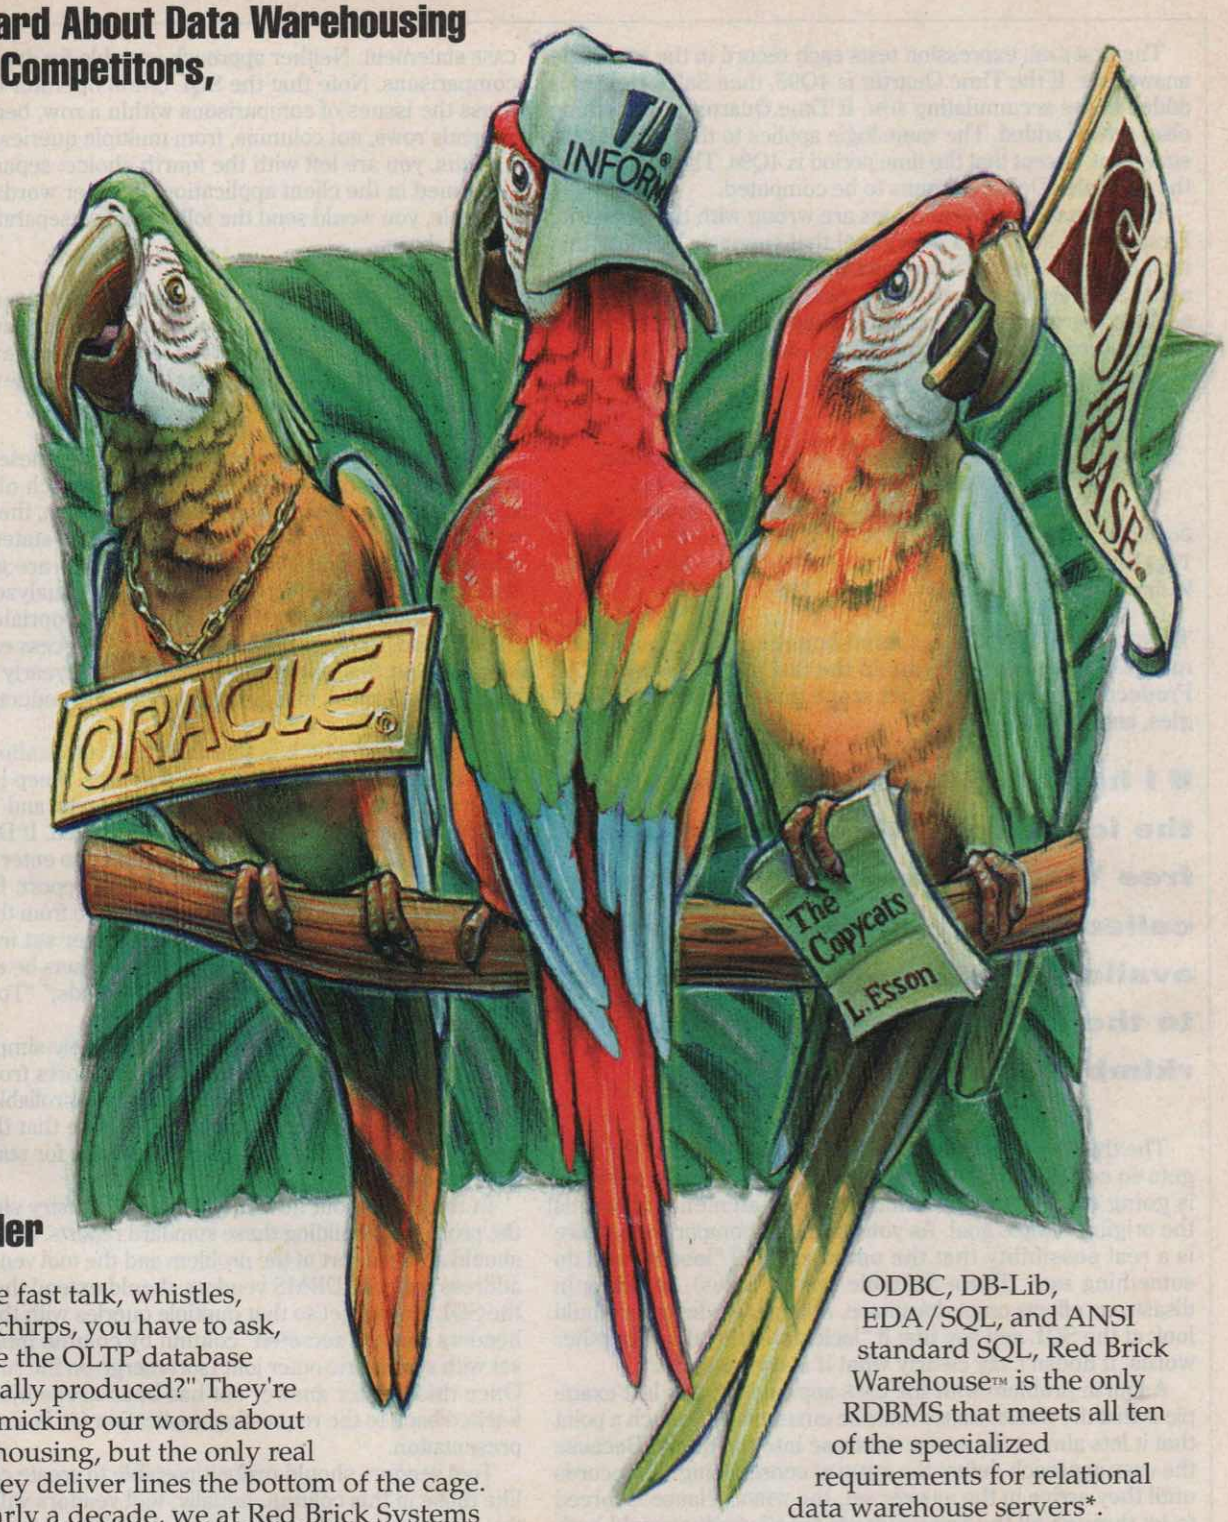
```
WHERE . . . Time.Quarter IN ('4Q95', '4Q94')
```

then the SUM expression adds up the sales from both quarters, which you do not want.

SQL-92 provided a CASE expression that would seem to offer a way out. Put both quarters' data into the query, and then sort out the two quarters in the SELECT list. This is a kind of IF-THEN logic:

```
SELECT Product.Product_Name,
  SUM(CASE(Time.Quarter = '4Q95', Sales.Dollars, 0)),
  SUM(CASE(Time.Quarter = '4Q94', Sales.Dollars, 0))
FROM Sales, Product, Time
WHERE . . . Time.Quarter IN ('4Q95', '4Q94')
AND Product.Product_Name in ('Doodads', 'Toggles',
 'Widgets')
GROUP BY Product.Product_Name
```

You've Heard About Data Warehousing From Our Competitors,



Now Hear About it From the Leader

Beyond the fast talk, whistles, snaps and chirps, you have to ask, "what have the OLTP database vendors really produced?" They're great at mimicking our words about data warehousing, but the only real product they deliver lines the bottom of the cage.

For nearly a decade, we at Red Brick Systems have pioneered the data warehouse industry and have built an open relational database just for data warehousing. Supporting all popular UNIX platforms, 60 tool partners, ODBC, DB-Lib, EDA/SQL, and ANSI standard SQL, Red Brick Warehouse™ is the only RDBMS that meets all ten of the specialized requirements for relational data warehouse servers*.

In this game of "follow the leader," the important thing to remember is who the leader really is: Red Brick Systems. Call us today—we'll tell you what we're going to teach those other birds to say next.

RED BRICK®

The Data Warehouse Company™

* for a free copy of the ten specialized requirements for relational data warehouse servers, call Red Brick today!

1-800-777-2585

Oracle®, Informix®, and Sybase® are registered trademarks of their respective owners.

CIRCLE READER SERVICE NUMBER 56

The first CASE expression tests each record in the candidate answer set. If the Time.Quarter is 4Q95, then Sales.Dollars is added to the accumulating SUM. If Time.Quarter is something else, zero is added. The same logic applies to the second CASE statement, except that the time period is 4Q94. This logic allows the two Sales.Dollar columns to be computed.

Unfortunately, several things are wrong with this scenario. First, I don't know of a query tool that supports building this kind of logic. The CASE construct would have to be hand-coded at the SQL level into applications. Second, the approach doesn't scale to more realistic examples. Suppose you wanted to do a more interesting set of comparisons, as follows:

	4Q95 Sales	4Q94 Sales	1995 Sales	1994 Sales	4Q95 Sales Fraction to All Products
Doodads	57	66	210	213	16%
Toggles	29	24	110	93	8%
Widgets	115	89	409	295	32%

This example turns into a mind-boggling set of CASE statements. Notice the little trap in the last column, where "All Products" represents a larger scope than just Doodads, Toggles, and Widgets.

■ I have implemented several of the ideas from this column in a free Visual Basic report writer called Star Tracker, which is available to anyone with access to the Internet (http://www.rkimball.com).

The third problem with the CASE approach is that the SQL gets so complicated that the optimizer can't really tell what is going on. The proliferation of CASE statements obscures the original simple goal. As you add more comparisons, there is a real possibility that the optimizer will "lose it" and do something zany (like a full table scan of Sales), resulting in disastrous effects on performance. A language designer would look at the SQL and say that it "lacks intentionality." In other words, it doesn't say clearly what it is trying to do.

A fourth problem with the CASE approach in this last example is that the WHERE clause must be emasculated to such a point that it lets almost the entire database into the query. Because the CASE approach defers the issue of constraining the records until they arrive in the answer set, the WHERE clause is forced to let through all the Sales records for all products sold both this year and the previous year. If you had hoped to use a speedy aggregate record (see November's column, "The Aggregate Navigator," page 18), you will be disappointed, because the use of the All Products aggregate is hopelessly embedded in the query using the CASE approach.

Stepping back from the details, how do you program comparisons using a relational database? Historically, you have had only four choices: CASE statements, SQL self-joins, SQL correlated subselects, and separate queries combined in the client application.

The second and third approaches (self-joins and correlated subselects) are older techniques that are even uglier than the CASE statement. Neither approach is viable for data warehouse comparisons. Note that the SQL UNION operator does not address the issues of comparisons within a row, because UNION appends rows, not columns, from multiple queries.

Thus, you are left with the fourth choice: separate queries combined in the client application. In other words, in the last example, you would send the following five separate queries to the database:

```
4Q95 sales for Doodads, Toggles, and Widgets,
4Q94 sales for Doodads, Toggles, and Widgets,
1995 sales for Doodads, Toggles, and Widgets,
1994 sales for Doodads, Toggles, and Widgets, and
4Q95 sales for All Products.
```

Your client application must then combine these results outside of the database. Although this approach obviously increases the complexity of the client application, there are some valuable benefits. First, each of the five SQL statements is extremely simple SQL (which all query tools are good at producing). Second, the optimizer can easily analyze these simple SQL statements in order to choose appropriate evaluation plans. Third, the aggregate navigator will process each of these separate queries smoothly, jumping to the yearly aggregates for queries 3 and 4, and jumping to the All Products aggregate for query 5.

The only real glitch is that the client application must perform an outer join on the five answer sets. Keep in mind that there is no guarantee that Doodads, Widgets, and Toggles are all sold in the various time periods requested. If Doodads had not been sold in 1994, you would be forced to enter null data elements into the appropriate cells of your report. Fortunately, outer join is just a fancy name for sort-merge from the old Cobol data processing days. Building the answer set from the five queries requires only that the five answer sets be sort-merged on the set of sorted row headers ("Doodads," "Toggles," and "Widgets").

The separate query approach is extremely simple and general. You can build very complicated reports from separate queries, and the database performance is controllable and scales linearly in the number of queries. I believe that the separate query approach is the only viable approach for standard business reports.

In thinking about how the database industry should attack the problem of building these standard reports, DBMS vendors should address part of the problem and the tool vendors should address the rest. DBMS vendors should extend the concept of the SQL answer set so that multiple queries with the same row headers can be "accreted" column by column into an answer set with symmetric outer join (sort-merge on the row headers). Once this master answer set has been assembled, it can be fetched back to the requesting application for final analysis and presentation.

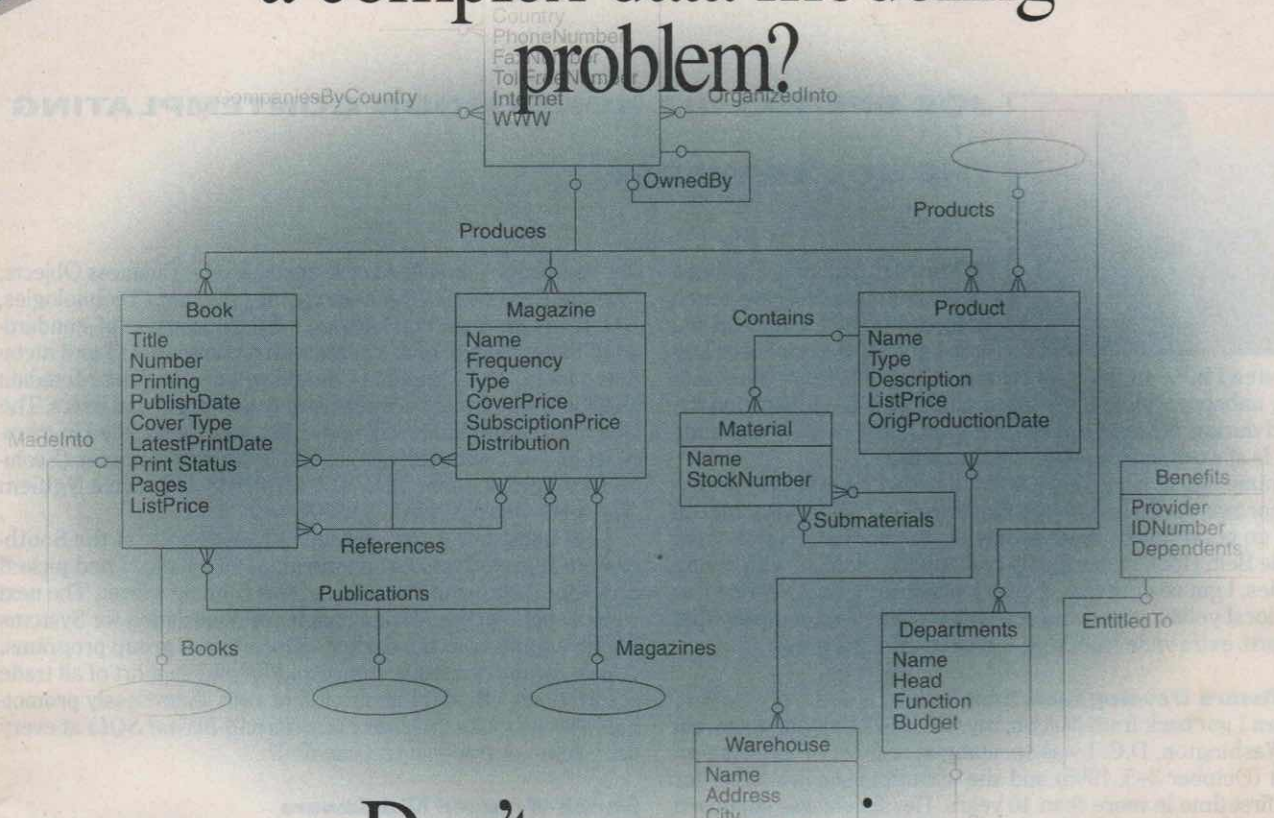
Tool vendors should make it possible to create comparisons like those in this column. Initially, tool vendors will have to do the sort-merge in the client; later, DBMS vendors can evolve this capability on the server side.

As an example (and as a goad to the industry), I have implemented several of these ideas in a free Visual Basic report writer called Star Tracker, which is available to anyone with access to the Internet (http://www.rkimball.com). This free software is also available on the *DBMS* Buyer's Guide CD-ROM, shipped with this issue of the magazine. Star Tracker works on Access, Red Brick, Oracle, Informix, and Sybase databases.

Next month I will conclude my analysis of "Business versus SQL" by looking at sequential processing, which is not just difficult, it's impossible. ■

Attention Database Developers:

Do you have a complex data modeling problem?



Don't compromise.

Solve it with

TITANIUM DATABASE ENGINE

Optimal Performance for Complex Development Needs

Other databases require design concessions when modeling complex data. The TITANIUM™ database engine directly supports highly complex relationships, without redundant data, without intersect tables, without joins... without compromise.

Parts explosion problems, for example, and other many-to-many, recursive, and n-ary modeling constructs are handled simply and elegantly by TITANIUM's patent-pending Dynamic Pointer Array (DPA) technology.

TITANIUM is a multi-threaded client/server and stand-alone database management system, featuring navigational API's for all major programming languages, SQL/ODBC support, compiled "C" stored procedures, fault-tolerance, transaction processing, and on-line backup — all contained within an extremely small memory footprint. Supported platforms include DOS, Windows (3.x, 95, NT), OS/2, NetWare, and UNIX.

For a TITANIUM Technical White Paper, call ***800-445-mdbs*** *or* ***317-463-7200****.*

E-Mail: info@mdbs.com WWW: http://www.mdbs.com

Successful Applications Are Built With TITANIUM.

mdbs®
Micro Data Base Systems, Inc.

mdbs is a registered trademark and TITANIUM is a trademark of Micro Data Base Systems, Inc. ©1995 Micro Data Base Systems, Inc.

CIRCLE READER SERVICE NUMBER 155

BY JOE CELKO

A Little R & R



JOE SPENDS HIS HONEYMOON CONTEMPLATING THE SQL UNIVERSE.

The SQL Explorer has been traveling a little too much lately and would love to stop and rest. I got married on Labor Day weekend in September, and immediately made my new wife very unhappy with me. My travel schedule is keeping me on the road during the week, and when I am home on the weekends she is at work, pulling two 12-hour shifts.

Immediately after the wedding, I had to go to Macon, Georgia for two weeks to look at a data warehousing project. Macon has no good hotels, and the town is the buckle of the Baptist Bible Belt. Hooters is a family restaurant, complete with crying babies. I got so desperate I even looked up "Escort Services" in the local yellow pages, but the only listing was a company that escorts extra-wide trucks on the interstate highways.

Software Development East '95

When I got back from Macon, my wife and I immediately went to Washington, D.C. I was speaking at Software Development East (October 3–5, 1995) and she wanted to see her father for the first time in more than 10 years. Her father was supposed to give us a ride from Dulles International Airport, but this didn't happen, so I missed my book signing at the Computer Literacy Bookstore at the show. (However, even though I wasn't there, the bookstore sold more than 40 copies of my book, *SQL for Smarties*.) Software Development is one of the last of the hard-core dweeb conferences, complete with "real" programmers and techies from all over the world. They talk shop, not management.

At the end of the trip, we connected up with my father-in-law and he managed to miss the Dulles International Airport turnoff three times when he gave us a ride back. This was our honeymoon.

DCI Data Warehousing Conference

When we got home from D.C., I had to take care of my CompuServe mail — during the month of October I ran a study group in the *DBMS* forum, based on my book. Then I got one day at the office to clear up 28 voice mails before heading out to Phoenix for three days for the DCI Data Warehousing Conference. The conference had approximately 900 attendees, more than double last year's attendance.

The META Group announced at the conference that it is involved with the newly formed Metadata Council. This is a vendor consortium composed of Arbor Software, Business Objects, Cognos, Evolutionary Technology Inc., Platinum Technologies, and Texas Instruments Software, with the purpose of standardizing the exchange of metadata with a common API and metadata model. The Metadata Council will oversee the Metadata Coalition, which will be made up of vendors and end users. The Council's membership will rotate. If you are interested in membership, you can contact either Dr. Katherine Hammer (Evolutionary Technologies Inc., 512-327-6994) or Patricia Nghiem (Business Objects, 408-973-9300).

I got back on Friday, October 13, and spoke at the Southeastern Ingres Users Group meeting. Of course, I had picked up a lung infection in Phoenix and was running a fever. The next week I spoke to the Atlanta chapter of Association for Systems Management. I am the darling of local users group programs, known for my generous nature and helpful support of all trade associations. Oh, did I mention that I am shamelessly promoting my two books (*SQL for Smarties* and *Instant SQL*) at every user group or trade show I can find?

Shrink-Wrapped Middleware

After I got back from the Data Warehousing Conference, I had enough time to wash my clothes, read my mail, and repack my bags for another trip. This time it was back to Minneapolis to do some follow up on Ensodex (Shoreview, Minn., 612-638-0721, http://www.ensodex.com); see my column, "The World of Software," June 1995) and its product. John Paulson and Harold Slawik picked me up at the airport and took me over to their downtown offices. Since my last visit, the name of the product has become HotSockets, and it is now ready to go to market.

Ensodex provides a set of drivers and middleware that lets you connect from any client to any server via ODBC. The client can be either Unix, Windows NT, Windows 95, or any product that supports an ODBC driver. While we were waiting for the graphic artist to show up with the final packaging from the printer, I got to play with the product using a PC-based copy of Microsoft Query. The connection was from the PC to an Oracle database running on a Sun machine at the Ensodex development office. I connected to the database (located somewhere in New York state) via IBMnet and the Internet, and then back to Minneapolis — not over local telephone lines. The response time was as good as I normally get from my admittedly old machine with a local database at home. The speed is limited by the modems, the quality of the lines, and the I/O time at the front end.

The bottom line is that Internet ODBC is now real and available at commodity prices. This means that it is practical to use the Internet as your company network. When we got back to the development office, I broke open a shrink-wrapped package (containing diskettes and a manual) and performed an installation first on Windows NT and then on Unix. I encountered no prob-

Joe Celko is a member of the ANSI X3H2 Database Standards Committee, a widely published author, and a consultant with O-Soft Development Corp. in Atlanta. Joe is also a frequent contributor to the DBMS *Forum on CompuServe. You can reach him at 71062,1056.*



Reach for the Stars!

All the solutions you ever wanted from an **Oracle** database utility in **one**, fully integrated, Windows product.

- ★ Performance Monitoring
- ★ Alarming & Alerting
- ★ Tablespace Management
- ★ Table/Tablespace Reorganization
- ★ Automatic Capacity Management
- ★ Trend Analysis
- ★ Database Administration
- ★ Change Management

For more information, or to receive your **Free** "QuickStart" kit...
Call Toll-Free: **1-800-275-2723**

Bradmark UK Ltd: Tel: (0) 1905.795444
Fax: (0) 1905.795531

Business Alliance Programme

SYBASE OPEN SOLUTIONS PARTNER

BRADMARK Technologies, Inc.

Providers of premier data management solutions.

DBGENERAL: The Database Console is a trademark of Bradmark Technologies, Inc. Other product names herein have been used for identification purposes only, and may be trademarks and/or registered trademarks of the respective companies.

CIRCLE READER SERVICE NUMBER 278

The First RAD Database Development Tool For C++

ProtoGen+ Client/Server Suite Version 5.0

The Fastest Way to Build High-Performance Desktop and Client/Server Applications

When it comes to visual programming, there is no equal. Imagine if you had the productivity of a 4GL with the power of C++. With ProtoGen+ you have it! You get a collection of tools designed to make creating screens that interact with your database simple, and you still have the C++ code that you can modify and extend.



With our visual programming techniques, (no coding) you can create SQL-queries, forms with one-to-many relationships, shared data sets across forms, graphs linked to ODBC data sources or to our DataTable grid control. Create database validation lists, define help links, tool tips and add data validation. Play sound clips, do drag and drop and create great looking reports.

But that's not all - We have Advanced Team Development.

More than just version control integration for file check-in, check-out and project synchronization. Our workbench manages teams of programmers building applications in parallel. Each developer works in ProtoGen+ to visually design screens, access databases and create other application components.

ProtoGen+ generates a DLL for each sub-unit with expert level C and C++ Code Generation for **Windows**, **Windows NT** and **Windows 95**.

Reusable Object-Oriented Code

Most of your application can be created using our visual programming techniques, but not all. Which is why we generate industry standard C and C++ code for the latest versions of MFC and OWL. You can choose to generate 16 bit Windows code or 32 bit Windows NT or Windows 95 code. Our regeneration technology will preserve your code when regenerating your source modules.

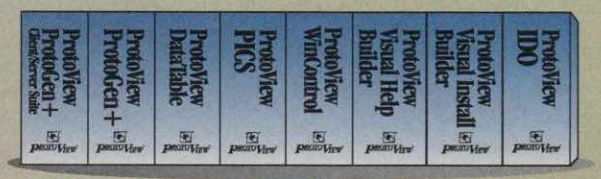
Test Mode for Interactive Development

ProtoGen+ also supports a test mode that's designed to help you do interactive development without constantly having to generate, compile and link your application. This unique feature also supports testing screen navigation, database access, data validation, graphs and reports.

If you don't think that's enough...

We include royalty free 16 and 32 bit Intersolv ODBC data drivers for both Desktop and SQL databases.

If you want to build high performance applications fast, Call now 1-800-231-8588.



The ProtoView family of software components and Visual Programming tools.

PROTOVIEW™

The Visual Development Edge™

Call Today 1-800-231-8588

Tel: (609) 655-5000 Fax: (609)655-5353 Internet: HTTP://WWW.PROTOVIEW.COM

Copyright © 1995 ProtoView Development Corporation,All rights reserved.All ProtoView product names are trademarks of ProtoView Development Corporation.

CIRCLE READER SERVICE NUMBER 15

lems, other than some forgotten Unix commands. You simply call a set-up program, then go to a GUI form screen to configure your data sources. HotSockets works with Intersolv, Visigenics, Windows NT, and Windows 95 ODBC drivers.

Also, HotSockets has a block-fetch feature that you can turn on and off during a session. The typical mode is to pass one result row at a time over the wire. However, with a block fetch, the result set is sent in large blocks. This is a faster approach, because it greatly reduces network traffic. The trade-off is that you cannot use SQLGetData( ), SQLSetPos( ), or positioned UPDATEs and DELETEs. In my June column, I mentioned that HotSockets was running at 200 rows per second. With block fetch, the same data (25,000 + rows of 70 characters) ran at better than 1300 rows per second.

The design philosophy is "save the server" — that is, if one client fails, the server remains running for the other clients. There are also error and communication logs, so you have some idea of what happened in the event of a failure.

One unexpected feature of the product is that you can use it as a gateway by connecting client A on machine A to a HotSockets driver on machine B, which connects to a HotSockets Driver on machine B that, in turn, connects to machine C. (Imagine extension cords strung together.) The user of machine A is aware only that the machine is running a little slower than it should. The user does not see any of the connections until an error message appears that shows the chain of connections between the client and the database on the back end.

I want to see what other people do with this product. It needs a security package for heavy commercial applications on the Internet, but these types of packages are available off the shelf. Data warehousing applications can get online access to production databases, and replication services are suddenly easy to set up using script languages that can handle ODBC. Right now, Ensodex is working on a data source name service that will provide clients with the names of their ODBC data sources on the Internet. Ideally, a company would distribute a diskette with HotSockets, and the installation would create a file of all the sources to which the users have access.

ANSI X3H2 Meeting in Jersey

I left Minneapolis for Long Branch, N.J. and another ANSI X3H2 Database Standards Committee meeting (October 23–26). Luckily, the Hampton Inn in Minneapolis had guest washing machines.

Long Branch is a town on the Jersey shore that has one really nice resort hotel, where we stayed, an excellent crab house restaurant, and nothing else. The town consists of houses and businesses that are boarded up and covered with For Sale signs. I was told that the beach area burned down in a storm a few years ago and was never built back because of political problems.

The X3H2 Committee is holding at 26 members (Borland dropped out and UniSQL returned in October). This meeting was only two days long, thanks to a short agenda.

Our beloved chairman, Don Deutsch, chided us for not making better progress on advancing SQL3. He stated his fear that it could become a moot point as the market overtakes the standards committees and provides de facto standards for object-oriented databases and the other areas we are trying to include in SQL3.

There is good evidence of this trend. The ANSI Z39 Standards Committee has been invited to speak to us about its Z39.50 text-retrieval language and other textbase problems. The Open Geographical Information Systems (OGIS) vendor consortium has overlapping membership with the ANSI X3L1 Committee, and it might define its own GIS query language shortly. The X/Open consortium is continuing work on its Call Level Interface (CLI), even as ANSI and ISO advance the SQL/CLI to international status (by the time you read this column in January, ISO could have already approved it).

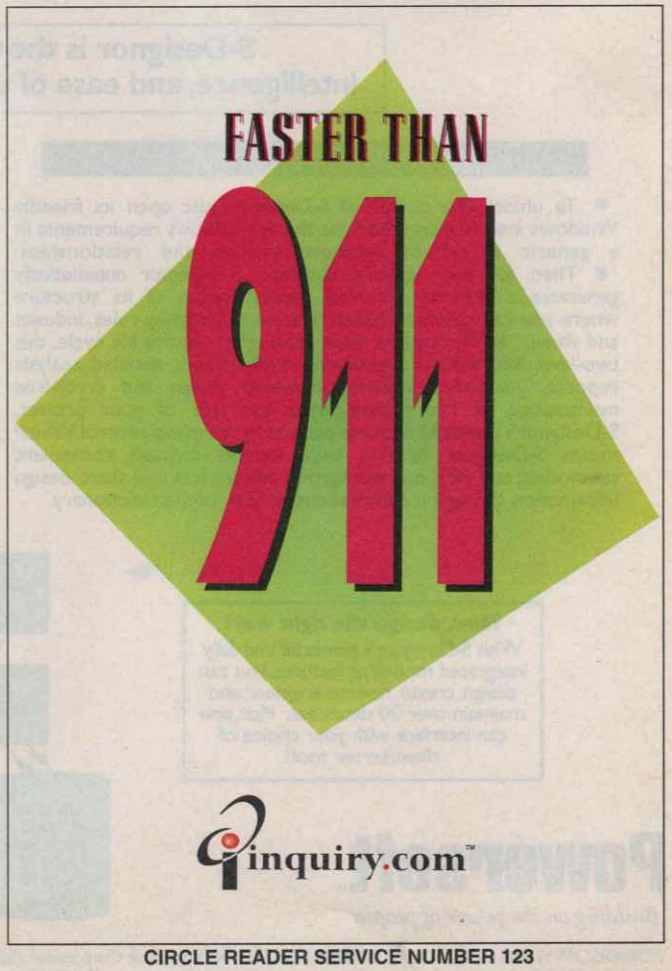
ANSI X3H2 makes good use of email to pass around documents and meeting information, but this is not the case in ISO. Instead, it makes a lot of money selling hard-copy standards documents, so you can understand its lack of desire to publish these documents on the Web. However, this fear extends to internal committee use, too. The July 1995 meeting of ISO/IEC JTC 1/SC24 approved a trial of electronic documents for the period of September 1995 to May 1996. The U.K. is against electronic documents, while Germany and Italy want to study the use of diskettes and "sneaker net" as a first step.

The one interesting new idea I got at the meeting was "logical locking"; that is, instead of putting physical locks on rows or pages of storage, you route all database requests to a referee who decides what you are allowed to update. Then, if one user wants to update all the rows where (age>21), and another user wants to update all the rows where (hair = 'Blonde'), the referee must determine if there is any overlap in the result sets, and exclude one or both transactions from the database. This is another way to free the logical model from the physical model.

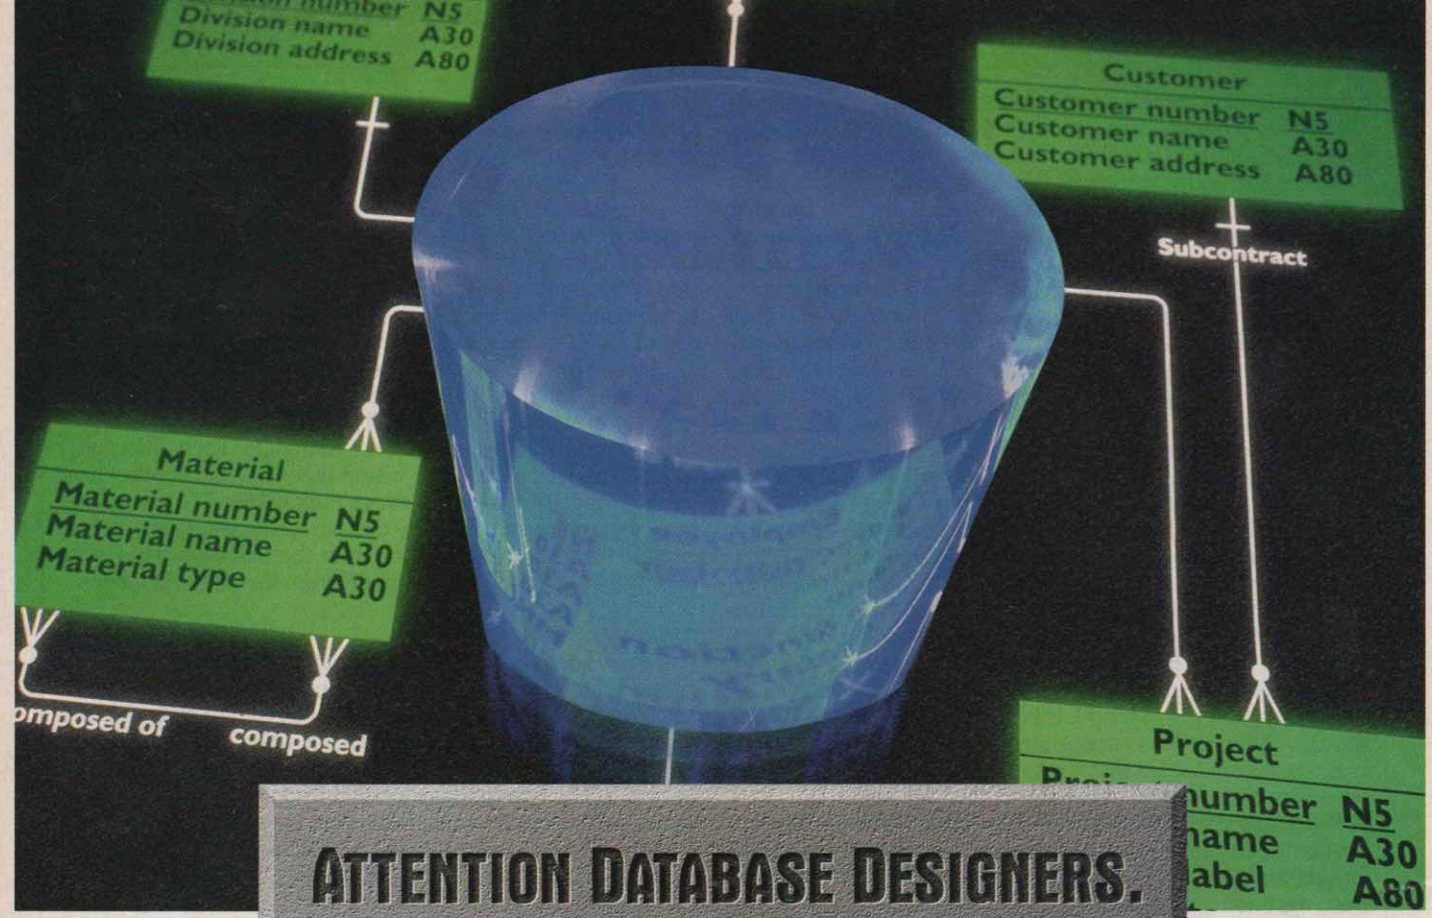
Sybase Makes Moves

Phil Shaw announced in an October email broadcast that he has changed his affiliation to Sybase (Emeryville, Calif.). Phil formerly worked for IBM as its point man on ANSI X3H2, and then moved over to do the same job at Oracle when IBM "downsized." Phil is one of the oldest and most technically competent members of X3H2. Sybase has already hired two former DEC representatives as well as Chairman Don Deutsch.

I find Sybase's move to buy up X3H2 people interesting because Sybase is at a critical point. On the technical side, its SQL Server product has been noted for poor scalability, bugs, and a non-standard, non-portable pseudo-SQL language. On the financial side, the company lost $17.4 million the first quarter of



CIRCLE READER SERVICE NUMBER 123



ATTENTION DATABASE DESIGNERS.

S-Designor is the tool with the power, intelligence, and ease of use you've been looking for.

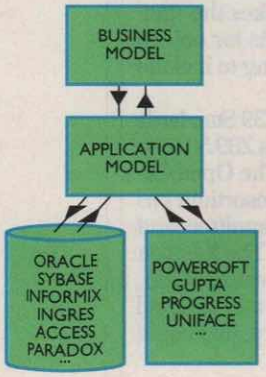
Superior database modelling.

■ To unlock the power of S-Designor, just open its friendly Windows interface and describe the application's requirements in a generic model of business entities and relationships. ■ Then, for your specific database, S-Designor *automatically* generates a separate graphical representation of its structure where you can optimize tables, references, integrity rules, indexes and views. ■ Throughout your application's entire life cycle, this two-level Information Engineering method, plus detailed analysis reports, guarantee superior database design and error-free maintenance. ■ No matter what the size of your project, S-Designor's powerful features put you in complete control. Which means S-Designor handles large models through convenient submodels; and with our workgroup edition, lets you share design information through a comprehensive SQL project dictionary.

Databases you can depend on.

■ Whether you want to create, reverse-engineer, maintain, upsize, downsize or rightsize your database, S-Designor is the smart choice. ■ And S-Designor does it right for over *thirty* databases including Oracle, Sybase, Informix, Ingres, SQL Server, Progress, DB2, AS400, SQLBase, Paradox, Access, FoxPro and many others. From the design models, S-Designor automatically creates and maintains the optimal database structure, including constraints, views, triggers and stored procedures. ■ For maximum client/server performance, S-Designor *fully bridges with PowerBuilder, Gupta SQL Windows* and other leading 4GL tools. ■ The bottom line? If you're looking for a database design tool with incredible power, intelligence, and ease of use, take a close look at S-Designor. It's the easy way to design truly perfect databases.

Now, design the *right* way!
With S-Designor's powerful and fully integrated modelling features, you can design, create, reverse-engineer and maintain over 30 databases. Plus, you can interface with your choice of client/server tool!



Free S-Designor evaluation kit.

With S-Designor it's actually easy to design perfect databases. But don't take our word for it; prove it to yourself. Call right now and we'll rush you a fully-functioning copy of S-Designor that lets you evaluate all its powerful features.



1-800-395-3525

Powersoft
Building on the power of people.

POWERSOFT IS THE OPEN TOOLS DIVISION OF SYBASE. Powersoft Corporation, (508)287-1500. http://www.powersoft.com/ Powersoft Europe Ltd., United Kingdom, +44 1494 555555. American Powersoft Corporation Pte Ltd., Singapore, +65 338-0018. All trademarks and registered trademarks are property of their respective owners.

CIRCLE READER SERVICE NUMBER 115

1995, and fell short of its 1994 earnings the rest of year.

To solve some of the technical problems, Sybase is pushing its new SQL Server 11 and a bitmap indexing method called Sybase IQ, and promises to ship products that support data warehousing. (See Client/Server Connection, page 88.) Sybase IQ is based on the Expressway technology, which Sybase acquired in early 1995. Other improvements should bring the company up to the performance levels that Informix and Oracle have already achieved in the last few years.

Sybase has also added better technical people. The purchase of Powersoft (and hence Watcom) gave the company some of the best compiler writers and SQL designers in the trade. Sybase is also buying up more individual X3H2 members than Computer Associates did when it was acquiring entire companies.

However, I am worried by Sybase's technical move to rename the next version of Watcom SQL (which should be version 5.0) to Sybase SQL Anywhere. This product will include a switch to run the old T-SQL language. Wouldn't the time have been better spent writing conversion tools to bring the victims of the old Sybase products up to current standards? How much do you have to cripple the perfectly good Watcom SQL engine to include mutant behavior? This is as dumb as producing an object-oriented autocoder to enable you to maintain backward compatibility.

If Sybase cannot get its act together, I will have to agree with Donald Feinberg of the Gartner Group, who said recently, "On a scale of one to 10, 10 being the ultimate RDBMS, Sybase needs a 20 here" (*Computerworld*, October 16, 1995, page 1). However, I think that right now Sybase could beat out or hold its own against Microsoft in this market niche. Microsoft suffers from being locked into NT as its only operating system at a time when more shops are looking at Unix systems. Microsoft's desktop database products, Access and FoxPro, are not compatible with SQL Server or with each other, they are not near any standard, and they are not designed for client/server applications. A reasonably priced suite of integrated tools from one vendor that can operate on multiple platforms would have a place here.

This Month's Puzzle

Mark Fronteraat of LanSoft Inc. (Miami) posted this month's problem in the CompuServe Watcom forum in September 1995. He has budgeting information that consists of the following three tables: the items to be paid for, the estimated amounts to be spent on them, and the actual amounts spent on them. Some items are covered by more than one check, and sometimes one check covers several items.

Items

```
item_no    description
==========================
10         'Item 10'
20         'Item 20'
30         'Item 30'
40         'Item 40'
50         'Item 50'
```

(continued on page 87)

PUZZLE ANSWER

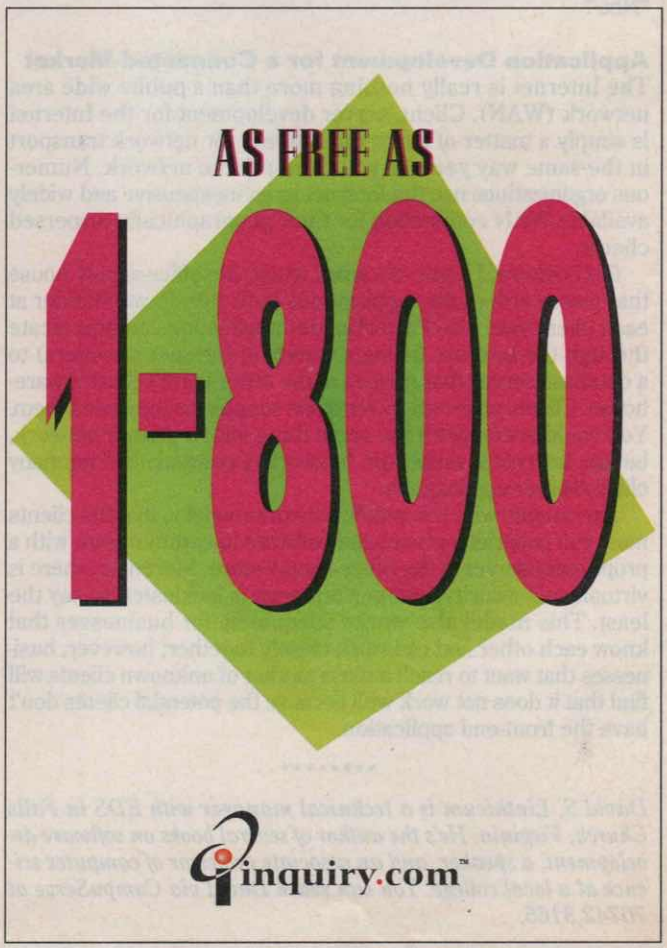
I think that this schema needs some work, but you can do this with scalar subqueries and some tricky code.

```
SELECT item_no, description,
(SELECT SUM (actual_amount)
FROM Actuals
WHERE Items.item_no = Actuals.item_no) AS tot_act,
(SELECT SUM (estimated_amount)
FROM Estimates
WHERE Items.item_no = Estimates.item_no) AS tot_est,
(SELECT CASE
WHEN COUNT (*) = 1
THEN MAX (check_no)
ELSE 'Mixed' END
FROM Actuals
WHERE Items.item_no = Actuals.item_no
GROUP BY item_no) AS check_no
FROM Items
WHERE tot_act IS NOT NULL
OR tot_est IS NOT NULL;
```

The trick is in the scalar subqueries. The first two subqueries calculate the total actual amounts and the total estimated amounts as if they were part of a GROUP BY and LEFT OUTER JOIN. I will offer a prize for a solution that uses a GROUP BY and LEFT OUTER JOIN and is still only one query (no VIEWS).

The final subquery is trickier. The query finds all of the Actuals that are associated with the item under consideration in the result table and makes a group from them. If the group is empty (no checks issued), the subquery returns a single NULL, and you display that NULL. If the group has one check in it, the CASE expression will return that single check number. The MAX( ) function is a safety check to guarantee that you have a scalar result from the subquery and you might not need it in all SQL-92 implementations. If there is more than one check actually issued on the item, then the COUNT(*) is greater than one and you get the string 'Mixed' instead of a string that represents the unique check number.

— Joe Celko



CIRCLE READER SERVICE NUMBER 123

By David S. Linthicum

The Client/Server Internet



WANT TO LEVERAGE THE INTERNET FOR CLIENT/SERVER? GRAB SOME HOTJAVA AND GET TO WORK.

The Internet is more than a passing technical fascination. It's an all-out explosion. It seems that everyone has Internet access in one form or another. Most (if not all) organizations larger than the corner hot dog stand have Web pages (including this magazine, at http://www.dbmsmag.com).

No, this is not another story about how the Internet is changing our lives. Instead, I look at how client/server developers can leverage the power of the Internet to build client/server applications for fun and profit. In particular, I explore how a new product from Sun Microsystems Computer Corp. (a unit of Sun Microsystems Inc., Mountain View, Calif.), called HotJava, could quickly become the best way to build client/server applications for the "Net."

Application Development for a Connected Market

The Internet is really nothing more than a public wide area network (WAN). Client/server development for the Internet is simply a matter of using the Internet for network transport in the same way you use your own private network. Numerous organizations use the Internet as an inexpensive and widely available WAN connection for their geographically dispersed clients.

For instance, I know of a small wholesale office-supply house that places order-entry applications built with PowerBuilder at each client site. The PowerBuilder applications communicate through the Internet (using a variety of Internet providers) to a database server that resides at the office-supply store's warehouse. Clients place orders for office supplies as they need them. You could accomplish the same thing with a private network, but the Internet is simply the "network of convenience" for many client/server applications.

The trouble with this public network model is that the clients must run proprietary front-end software to communicate with a proprietary server at the office-supply store. Moreover, there is virtually no security, and performance is lackluster, to say the least. This model also works adequately for businesses that know each other and can work closely together; however, businesses that want to reach a mass market of unknown clients will find that it does not work well because the potential clients don't have the front-end application.

.....

David S. Linthicum is a technical manager with EDS in Falls Church, Virginia. He's the author of several books on software development, a speaker, and an associate professor of computer science at a local college. You can reach David via CompuServe at 70742,3165.

Web browsers, such as Mosaic and Netscape, address this problem by enabling Internet users to access standard Hypertext Markup Language (HTML) documents from remote servers using standard transport protocols such as Hypertext Transmission Protocol (HTTP). There are thousands of businesses that have set up shop on the Internet using Web servers to communicate with the masses, including many that use HTML to communicate with database servers. (See "Database and the Internet," *DBMS*, December 1995, page 44.)

Using traditional Web browsers, users can view text and static pictures using HTML, as well as enter and look up data on more sophisticated Web servers using interactive Web server extensions such as CGI (Common Gateway Interface) protocol. However, the capabilities of HTML and CGI stop far short of what you would typically consider a fully functional client/server application. Client/server developers need a standard way to transport and execute free-form, platform-independent applications over the Internet, with the same ease and availability found in the traditional Web browsers.

Better Browsing with HotJava

HotJava (still in the alpha test stages at press time) builds on the Internet browsing technology of Mosaic and Netscape, and expands on their existing capabilities. Simply put, HotJava is more flexible than HTML. Developers can build Web applications with all of the functionality that they would expect from a traditional programming language such as C++.

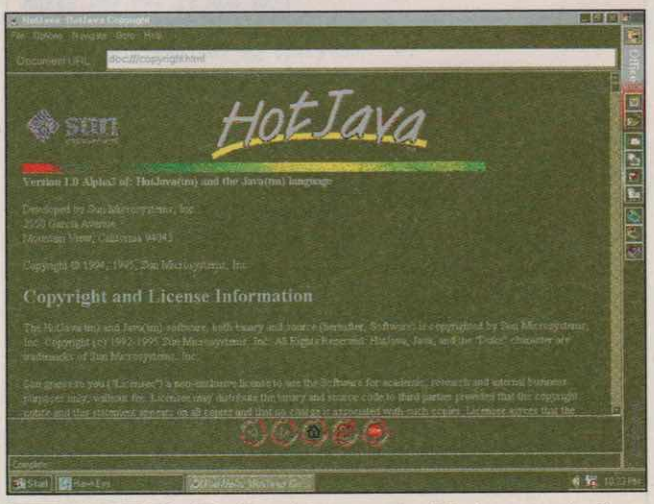


FIGURE 1 **The HotJava Browser is able to execute Java applets found on the Internet.**

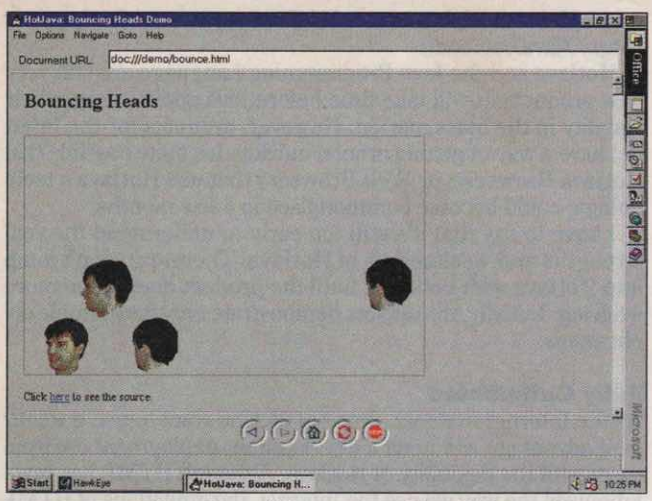


FIGURE 2 **The Java "Bouncing Head Demo" applet demonstrates the flexibility of the Java language as heads bounce all over the screen. The HotJava browser does not have to maintain a connection with a remote server when running applets.**

The HotJava Browser (see Figure 1, page 26) is very similar to traditional Web browsers, but HotJava can execute "applets." Applets are platform-independent programs written in the Java Programming Language and embedded, like images, in HTML documents.

HotJava lets users access applications over the Internet in a whole new way. One of the most exciting features of HotJava is that the software moves transparently over the network from a remote server, with no installation required. When you need to run a HotJava applet, HotJava downloads it from a remote server (along with any other software your applet may require) and invokes it on your local workstation. For example, Figure 2 depicts a demonstration applet program. Here, the heads bounce quickly inside of an applet window running inside of the HotJava browser. Applets can execute without a continuous connection to a remote server on the Internet. And because Java is completely open, developers can build applets that can handle streaming data from the server, including video broadcasts, stock feeds, sound such as Real Audio, and so on. Developers can define the application, as well as the way the application communicates with other resources such as database servers, video servers, and audio servers.

The use of the HotJava program execution model means that content developers for the World Wide Web don't need to concern themselves with the fact that users may not have installed and configured a special piece of software required to view data. With HotJava, it happens automatically, as long as the user is running a browser that supports HotJava, or running the HotJava browser. This is browser-dependent technology. This transparent acquisition of applications lets developers move away from the limitation of the fixed media types such as images and text, and allows them to do anything they want, such as access information on a remote database server.

A New World

The basic difference between HotJava and traditional Web browsers is flexibility. Although other browsers have a lot of detailed, static knowledge of how to handle several different data types, protocols, and capabilities required to navigate the World Wide Web, HotJava is able to extend these capabilities through the use of applets. HotJava provides an open and extensible facility that lets developers add new capabilities to browsers on the fly.

HotJava applets provide this flexibility through interactive content, dynamic types, and dynamic protocols. Interactive con-

MORE COMPLETE THAN
411

Switch from Sybase® to Borland™ to Oracle® to Microsoft® at
www.inquiry.com

You're in the market for a new application development tool, database, report writer or other software development tool.

Quick. Where do you go? Browse a couple of software magazines. Ask some developers. Read a few reviews. Make a short list and call for some brochures and demo disks.

The trouble is, this takes a lot of time. And how do you know when you've covered everyone?

Now you can access the complete universe of software development products from one web site, www.inquiry.com. Every important technical review. All the manufacturer's brochures, white papers, specs. You can download a demo, sign up for a seminar, or ask for a sales presentation—all on-line. Now it's all in one place. It's all free. And it's all at www.inquiry.com.

inquiry.com™

©1995 inquiry.com. 2655 Campus Drive, Suite 120, San Mateo, CA 94403. Telephone (415) 655-2828. Fax (415) 655-2820. Email info@inquiry.com. All rights reserved.
The inquiry.com logo is a trademark of inquiry.com. All other trademarks and registered trademarks are the property of their respective owners.

CIRCLE READER SERVICE NUMBER 123

tent means that developers can create a Java application using the Java Programming Language, with the guarantee that others running HotJava can execute the application on any platform that supports HotJava. For example, a wayward Java applet cannot reformat your hard disk because Hot Java, the engine that runs the applets, won't allow it. This is very similar to the Telescript remote programming tool I wrote about last year ("A Better RPC," *DBMS*, June 1994, page 24). The Telescript agent had to have a local engine in order to run. The engine would not let the agent perform anything that would jeopardize the health of the local system. The same goes for applets. In addition, users know that HotJava applications contain no malicious code that can damage their systems. HotJava is ready for these sorts of tricks, and code that attempts to be malicious or just plain buggy can't get around the firewalls that HotJava places around its applets.

The use of dynamic types means that HotJava understands different types of objects immediately. For example, if someone creates a new image format, the developer simply places a copy of the Java code for reading the new image format on the server with the images. Other Web browsers would have to be updated to read the new image type, or users would have to download and configure software to read the new images. In contrast, HotJava upgrades itself on the fly to adapt to different data types. (The upgraded applet exists only in local memory while the applet is running. You can, however, save the applet to your local drive for future use. Most users will run an applet as it exists on a Web server, in the same way that they view an image or run a video clip.) The applications for this are endless, including reading various database formats, or connecting to a database server over the Internet.

Dynamic protocols mean that HotJava can incorporate new communications protocols dynamically to meet the needs of a Java application. This capability has the most significance for the client/server developer, because HotJava clients can dynamically receive a protocol and middleware information in support of an application. This enables HotJava applications, for example, to link to remote database servers.

A user running HotJava needs only invoke a Java client/server application, which in turn moves and executes the application code automatically, and also installs the custom protocol and middleware layers required for server connections. Many protocols and middleware layers may exist at one time.

Java Programming Language

The engine driving HotJava is the Java Programming Language, which is very similar to C++. Sun was planning on using C++ as the applet development language of choice, but it ran into some limitations with C++, and thus created a custom version designed for HotJava applet development. Sun describes it as a "simple, object-oriented, distributed, interpreted, robust, secure, architecture-neutral, portable, high-performance, multithreaded, and dynamic language." From what I've seen, Java is an object-oriented programming language optimized for distributed heterogeneous applications. The applets are compiled into machine-independent bytecodes (similar to p-code), so that applets can move over the Internet and execute (using an interpreter) on any platform running HotJava.

The bytecode format (created using a compiler) is compact, which allows bytecode to move quickly over the Internet, and the Java Language compiler generates bytecodes for developers. Bytecodes are architecturally neutral, and therefore can run (using an interpreter) on a variety of CPU types, operating systems, and GUIs. You need only a Java runtime (part of HotJava) to interpret the bytecodes into native calls.

How do you begin to create Java applets? The HotJava Browser comes with complete Java Language API documentation. Sun is very eager to spread the word of Java, and provides a complete tutorial to create HotJava applets on the Internet (http://java.sun.com).

HotJava and the Java Programming Language are still very new products. It will take time before Java applets prove their viability in the mass market. However, products for the Internet have a way of getting around quickly. It's quite possible that HotJava Browsers, or Web Browsers that use HotJava's technology, could become commonplace in a few months.

I have to say that it's still too early to understand the real strengths and weaknesses of HotJava. Therefore, don't jump into HotJava with both feet until the product does some more evolving. Initially, the applets demonstrate only very simple applications.

Fully Caffeinated

As the Internet matures, those of us who want to use it to our best advantage will need a development deployment environment with the flexibility of HotJava. Although HTML has revolutionized the way we use the Internet, the limitations of HTML won't allow it to meet all our application development requirements. It's still very much a text and image system, and it was not really intended to be a true application development mechanism, even though many people are using it that way.

HotJava builds on the capabilities of HTML, allowing innovative developers to create platform-independent client/server applications for mass distribution over the Internet. As the Internet explosion continues, products such as HotJava will allow client/server developers to leverage the far-reaching power of the Internet. If you want to use the Internet to your best advantage, keep a close eye on HotJava. ■

• Sun Microsystems Computer Co. (a unit of Sun Microsystems Inc.), 2550 Garcia Ave., Mountain View, CA 94043-1100; 800-821-4643 (in California, 800-821-4642), 415-960-1300, or fax 415-969-9131.

WHAT IF you could offer a simple, inexpensive solution that provides hypertext-style access to **relational databases** on networks ranging in size from **LANs** to the **Internet**?

WHAT IF the data could reside on **multiple platforms** and be accessed by **Windows™**, **Macintosh™** and **UNIX** computers?

WHAT IF you didn't have to deploy **custom client software** to access the data?

WHAT IF you could do all this in an afternoon?

over tea.



Database Connectivity for the World Wide Web

http://www.aspectse.com/product/dbweb/dbweb.html

$495
per server

Aspect Software Engineering 808.539.3782 email: info@aspectse.com

CIRCLE READER SERVICE NUMBER 55



Works with Visual Basic 4.0 and Windows '95!

Get your Visual Basic applications off to a flying start.



Introducing ERwin for Visual Basic.

Go on an important client/server mission without a database design and you may never get off the ground.

ERwin for Visual Basic lets you integrate award-winning database design into your Visual Basic 4.0 environment,



▶ *Team development: Maintain and reuse Visual Basic validation rules, controls and properties in the ERwin dictionary, keeping your team flying in formation.*

providing the database control you need to develop at mach speed.

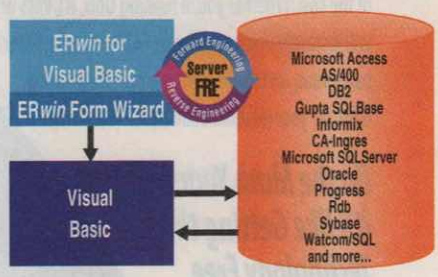
An ERwin database design gives your database the integrity and performance client/server applications demand.

And talk about productivity! ERwin can generate your SQL or Access database right from a data model, including triggers and stored procedures. Or reverse-engineer an existing database into a data model, ready for a new Visual Basic application.

As you prototype a Visual Basic application, ERwin lets you "see" its underlying database structure. This keeps your VB Forms and database in synch—as they evolve—without endlessly checking back and forth.

You can even specify Visual Basic validation rules, controls and properties within your database design and let ERwin's Form Wizard instantly generate "data aware" ready-to-run forms.

Don't fly blind! Using ERwin for Visual Basic, you'll soar through client/server missions faster and easier than you ever thought possible. With many happy landings.



▶ *Model-driven development: ERwin helps you understand and control your database as you build Visual Basic applications.*

Special Introductory Price! Call 1-800-78-ERWIN, dept. VB today or visit us at http://www.logicworks.com.

logic works®

Making Database Design Easier in a Client/Server World.™

Microsoft® SOLUTION PROVIDER

Logic Works and ERwin are U.S. registered trademarks of Logic Works, Inc. Logic Works with logo is a trademark of Logic Works, Inc. Visual Basic is a trademark of Microsoft Corporation. All other products mentioned are trademarks of their respective companies.

CIRCLE READER SERVICE NUMBER 200

Boost your Database Performance up to 600%. FREE.

Zitel SCP™ Storage Co-Processor technology is a unique combination of on-board cache memory and high-performance mass storage that delivers near-SSD (Solid State Disk) performance at a fraction of the cost. Whether you're running Unix, NT, VMS or Novell, SCP will instantly improve I/O and OLTP response, reduce batch runs, and lower transaction and operating costs for your relational database.

Do the Math Yourself: it's like Getting the Technology Free.

Zitel customers tell us time savings alone can repay the SCP purchase price in as little as a single day, but SCP technology goes on to save you money in other critical areas as well. Like the additional disk capacity SCP frees up at no extra cost to you. The CPU cycles gained. The ability to manage larger databases with smaller cache memory requirements. Reducing the number of drives typically required for RDB performance tuning.

And with SCP's convenient "plug and play" SCSI interface and "expert" caching technology, the need for multi-drive tuning is virtually eliminated.



SCP Performance: available in Tower, Rack Mount, Desktop or 1.5-meter Cabinet configurations.

Proven Technology from an Industry Leader...Yours FREE for 30 days.

As a world leader in high-performance storage technology, Zitel has more than a decade of experience with more than 3,000 units installed and supported worldwide in mission-critical environments.

In fact, we're so convinced SCP will dramatically improve your database server performance that for qualified buyers, it's yours free to test for 30 days— Just call Zitel and ask about our SCP Product Trial Program at **1-800/622-5020**.

ORACLE Business Alliance Programme℠

SYBASE® OPEN SOLUTIONS™ PARTNER

zitel®

MEMORIES OF THE FUTURE™

Zitel Corporation
47211 Bayside Parkway
Fremont, CA 94538
Tel: 510-440-9600
Fax: 510-440-9696

© 1995 Zitel Corporation. Zitel and SCP are registered trademarks of Zitel Corporation. All other product names and brand names are trademarks of their respective owners. All rights reserved.

CIRCLE READER SERVICE NUMBER 61

BY CLARA H. PARKES

HANDS ON

Magic 6.0

- Magic Software Enterprises Inc., 1200 Main St., Irvine, CA 92714; 714-250-1718 or fax 714-250-7404.
- Pricing: $2500 for a single developer.
- Minimum requirements: 8MB of disk space, 4MB of RAM, and 386 or later.

Magic has always been a strong development environment. But, because previous versions of Magic lacked support for Windows-based application construction, many client/server developers have passed Magic by. Today, with the release of Magic 6.0, Magic Software now offers true Windows-based "no code" development for client/server. Magic also provides multiplatform support and supports application distribution through partitioning. Let's see what new tricks Magic 6.0 for Windows can pull out if its hat.

The most significant feature of Magic is its lack of a formal programming language. Magic uses "table-driven" development, in which developers build applications visually by defining several tables from within the tool. With its new Windows support, Magic 6.0 provides a new Form Editor, better application migration mechanisms, a character-based form generation object, a new help facility, and native DDE support for sharing information with other Windows applications.

Magic is a true multiplatform tool, with support for Windows, DOS, NetWare NLMs, SCO Unix, OpenVMS, OSF/1, AS/400, AIX, SunOS, Solaris, HP-UX, DG/UX, CTOS, and Olivetti Unix. Database support is extensive as well, with native drivers for Oracle, Sybase, Informix, Rdb, CA-Ingres, DB2/6000, DB2/400, Btrieve, Xbase, C-ISAM, RMS, c-tree, and MF-Cobol. In addition, Magic provides network support for TCP/Pathworks, NetStar/Plus TCP, PC/TCP, PathWay/TCP, Winsock, and LAN WorkPlace/TCP.

Where's the Code?

Building an application using Magic involves three steps. First, specify the desired file from Magic's File Dictionary. Second, define fields and special processes for your application. Finally, build the user interface. Magic developers design the low-level details of an application by setting values in tables instead of using a proprietary 4GL.

The Magic File Dictionary stores information about an application's data and structure. Using this facility, developers can specify fields, keys, type, and other characteristics held in the Magic repository. In addition, the repository knows where the database exists and the type of database server the developers are using. As an added bonus, Magic defines fields automatically by linking to an existing database schema.

Client/server applications do not live on database elements alone, and Magic developers must define programs to set up application behavior. Developers do this by defining the basic program parameters, setting up the program execution flow, designing the GUI, checking the program for logic errors, and establishing the end-user menu system. This is accomplished with Magic's complex series of menus and dialog boxes, such as the Program Dictionary table (see Figure 1), rather than with free-form programming.

The Program (task lists) Dictionary is a listing of all the programs that support a Magic application. This is where developers link behavior and interfaces to data; developers use this facility to create and edit Magic's codeless programs.

Each program in Magic's Program Dictionary has a Program Tree. (See Figure 2, page 32.) A Program Tree lets developers create and edit a program structure. The root box, which represents the root task in a program, is created automatically. From here you can "Zoom" from a task box into the Program Tree to the tables used to define the task's components.

Magic programs may be made up of several subtasks that exist as children of the parent task (or root task), as depicted in the hierarchy of the Program Tree. You define tasks using Magic's Task Properties dialog, which you access by executing Magic's Zoom menu selection from the Program Tree. Here you can specify such things as task type, default key, default expression, row limits, and selection tables.

Painting a Picture

To define the interface, Magic provides a GUI forms painter. (See Figure 3, page 32.) At first glance, it looks like there is only a limited number of controls on the palette. However, Magic lets you redefine controls so they have a new look and feel for your application. All of this control customization is set through tables, not programming code.

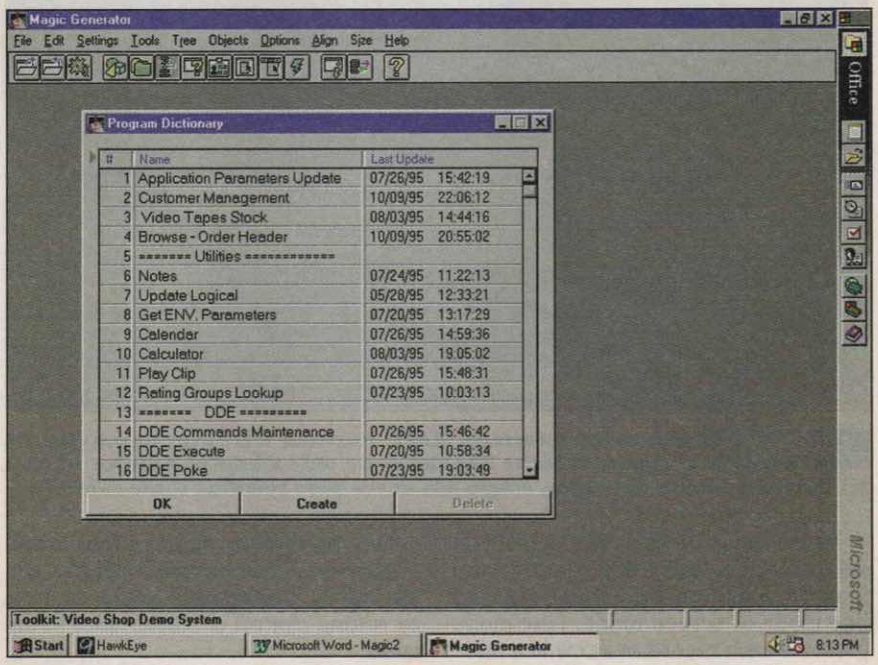


FIGURE 1 **Magic developers set up application behavior with Magic's complex series of menus and dialog boxes, such as the Program Dictionary table, rather than with free-form programming.**

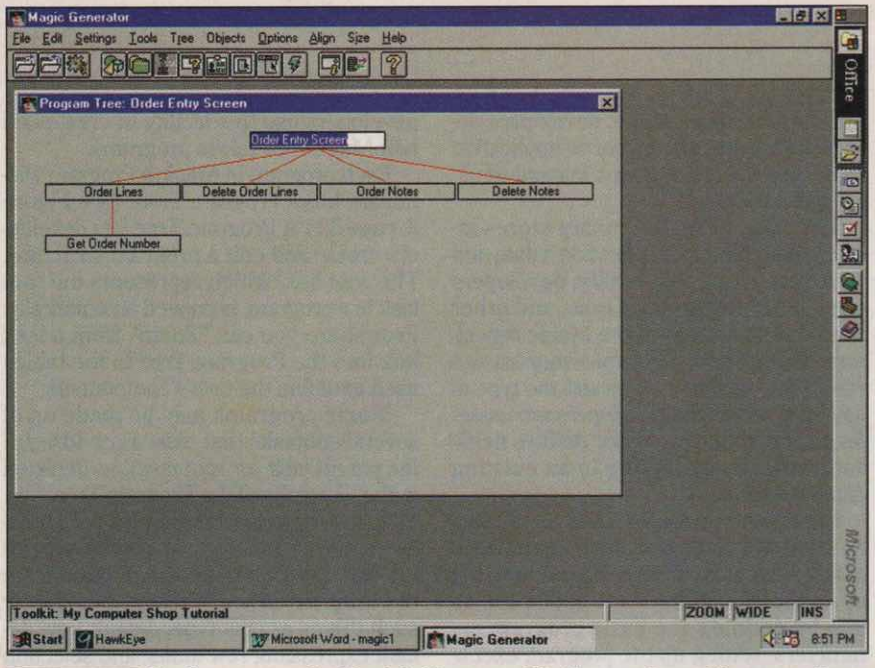


FIGURE 2 **Each program in Magic's Program Dictionary has a Program Tree. A Program Tree lets developers create and edit a program structure.**

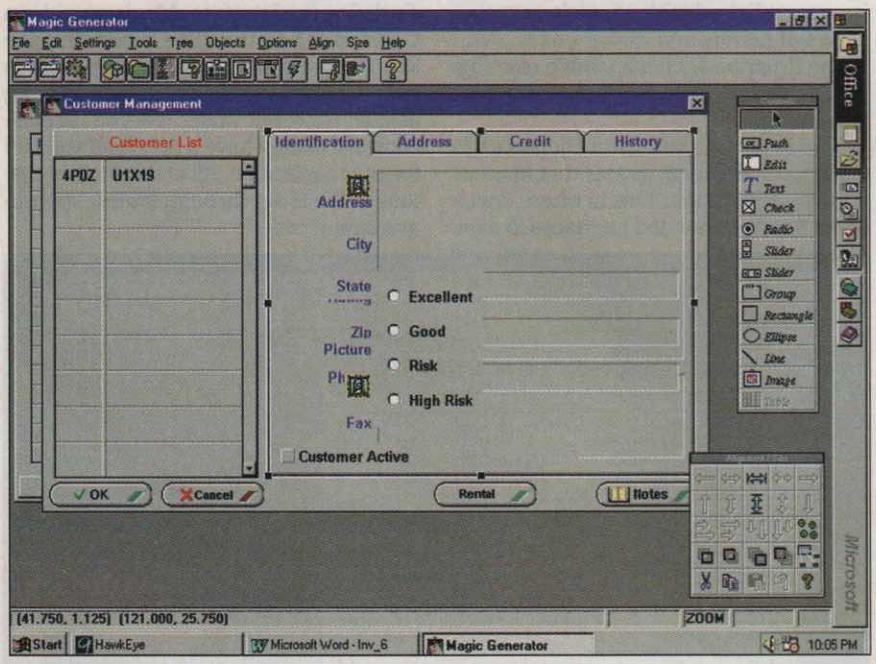


FIGURE 3 **To define the interface, Magic provides a GUI forms painter. At first glance, it looks like there are only a limited number of controls on the palette. However, Magic lets developers redefine these controls.**

Magic can generate a form when you select Options and Generate from the Magic menu. From here, you can select line mode or screen mode. Line mode presents each record as a row in a table format, whereas screen mode devotes a screen to each record. Once Magic generates a screen, you can modify the screen by clicking on the controls and placing them wherever you want. You can also define the properties of the controls by setting control characteristics in the table.

Once you define your Magic program, you can check your application to ensure that you made no logic errors by selecting Check Syntax from the Options menu. This is where Magic cruises through your program to find and report logic problems or table settings that don't make sense to Magic. Because no programming is involved, you can't make a syntax error. However, you must solve all the logic errors before your program can execute correctly. Because there is no code, you don't need a debugger. You can execute a new Magic program as you create it by selecting Execute Program from the Options menu. Performance in this interpreted environment is good on a 100 MHz Pentium.

Magic Partitioning

Like other client/server tools trying to gain a larger multitier market share, Magic supports application partitioning. Developers can partition Magic applications down to the task level. However, to execute tasks remotely on another server, you must have a native Magic engine running on that platform. For example, if you have a Magic engine running on an AS/400 or RS/6000, you can tell Magic to execute tasks there by simply providing the server names with the tasks from within the Magic environment.

This application partitioning scheme rivals that of application partitioning tools such as Dynasty and Forté, at only a fraction of the cost. Those tools, however, are built from the ground up for partitioning, and allow developers to partition and repartition applications with drag-and-drop simplicity. Magic still uses a manual table-driven process. Moreover, Magic still employs interpreter engines in lieu of native code generators and compilers. Therefore, performance is a bit lackluster in comparison.

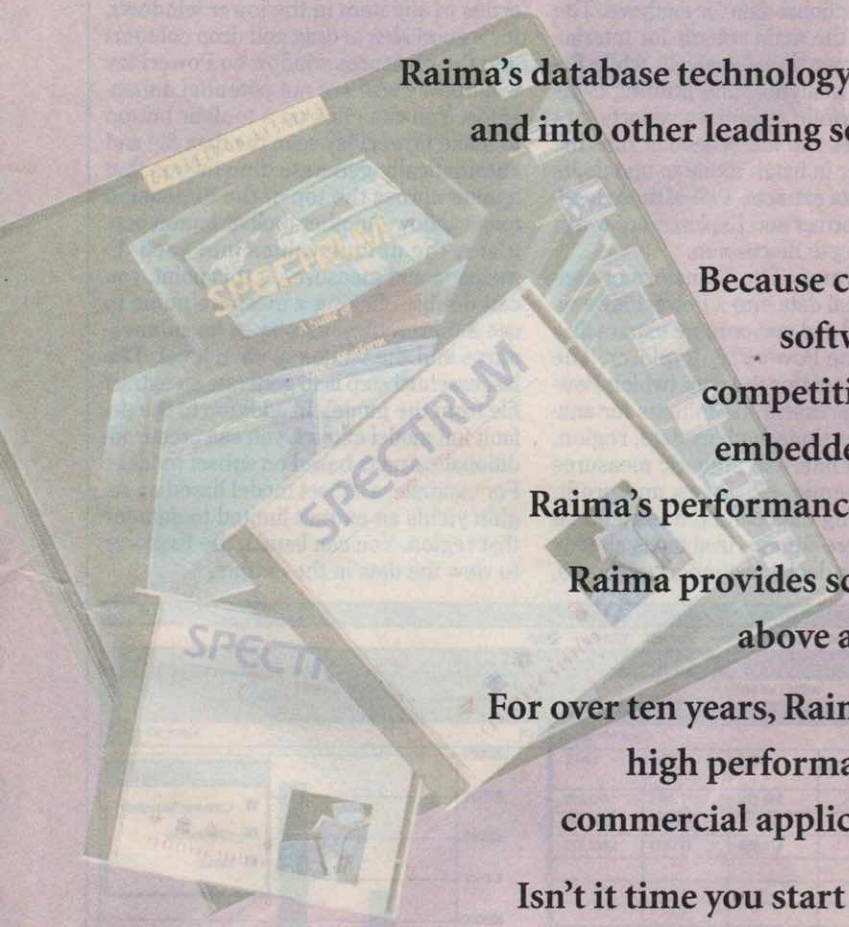
Your Own Magic Show

If you're considering Magic for your next client/server development project, you should know that it's complex, and it will take you some time to get up to speed with the tool. The documentation could also do a better job of explaining things. I found the tutorial confusing. Magic Software forces you to go through an annoying licensing procedure in which you must obtain an activation key from Magic Software to make the software fully operational. Magic does, however, give you a few tries with the software before shutting things down until you obtain your key.

Although the Magic development environment is "object-oriented-like" through its Program Dictionary, tasks, and subtasks, it does not fully support the object-oriented (OO) features you would find in other tools such as Delphi, PowerBuilder,

Corporations depend on Cabletron for network management.

Cabletron depends on Raima for database management.



Raima's database technology is built into Cabletron's Spectrum, and into other leading software products around the world.

Why?

Because companies like Cabletron want their software products to stand out from the competition. And because they've compared embedded database engines and have found Raima's performance and reliability are second to none.

Raima provides scalability, flexibility, reliability, and, above all, superior price and performance.

For over ten years, Raima has been providing embeddable, high performance database technology to leading commercial application developers around the world.

Isn't it time you start trusting the leaders and try Raima?

Raima Database Manager™
Raima Object Manager™
Velocis™ Database Server

Available for:	**DOS**	**HP-UX**
Netware 3&4	**OS/2**	**Solaris**
Windows NT	**SCO**	**QNX**
Windows 95	**AIX**	**MIPS ABI**

Call 1-800-ASK-4-RAIMA

Contact Raima for profiles of other demanding Raima customers.

RAIMA BASED

RAIMA™

If it's fast - it's Raima.

Raima Corporation, 1605 NW Sammamish Rd. #200, Issaquah, WA 98027 USA
Inside U.S.: 1-800-275-4724, Outside U.S.: 206-557-0200, Fax: 206-557-5200, Internet: sales@raima.com

Germany: 49 7022 9256 0 Singapore: 65 334 0061 Italy: 39 49 80 77 140 Argentina: 54 1 954 1414 Spain: 34 3 225 39 95
Raima UK Ltd.: 44 1 273 819 292 France: 33 1 30 67 29 63 Australia: 61 2 419 7177 Denmark: 45 44 88 99 00
Raima Benelux: 31 35 694 4738 Colombia: 57 1 62 10 070 Taiwan: 886 2 503 3006 Norway: 47 22 110 950 Sweden: 46 13 111 588
Baltic States: 372 631 4009 Finland: 358 0 8045 130 Japan: 81 3 5548 6926 Russia: 7 812 316 1965

CIRCLE READER SERVICE NUMBER 38

and SQLWindows. The non-OO table-driven method of development works well and is speedy, but it takes time to get used to it. I would also like to see support for VBX and OCX controls, in addition to the excellent DDE support that Magic already offers. I would also like to see a built-in database design facility. For now, you'll have to rely on another tool for that.

Putting all of the small problems aside, Magic is an excellent client/server development environment. You can develop both simple and complex applications in record time. This feature makes Magic an excellent candidate for a rapid application development (RAD) tool or prototyping tool. To top it all off, Magic is ready to accept your application on other environments that Magic supports without forcing you to modify your application. If you want to create powerful applications faster than the other tools available, you're going to need a little Magic.

.....

David S. Linthicum is a technical manager with EDS in Falls Church, Virginia. He's the author of several books on software development, a speaker, and an associate professor of computer science at a local college. You can reach David via CompuServe at 70742,3165.

HANDS ON



PowerPlay 4.1

- Cognos Inc., 67 S. Bedford St., Burlington, MA 01803; 800-426-4667, 617-229-6600, or fax 617-229-9844.
- Pricing: Single copy, $695 (includes PowerPlay Transformer, PowerPlay Explorer and Reporter user components, and PowerPlay portfolio).
- System requirements: Windows 3.1, 386 or later (486 is recommended), and 4MB of RAM.

Impromptu 3.0

- Cognos Inc., 67 S. Bedford St., Burlington, MA 01803; 800-426-4667, 617-229-6600, or fax 617-229-9844.
- Pricing: User version, $495; administrator version, $595.
- System requirements: Windows 3.1, 386 or later (486 is recommended) and 6MB of RAM.

One of the most populated segments of the client/service landscape is the end-user query, analysis, and reporting tools category. Cognos offers two fine products targeted at non-technical end users. PowerPlay is an analysis tool that emphasizes crosstabs and graphics. Impromptu is a report writer based primarily on prebuilt templates. Both require administrators to perform some set-up duties prior to unleashing them on users.

PowerPlay

PowerPlay consists of five modules: Transformer, Explorer, Reporter, Portfolio, and Scheduler. The Transformer prepares transactional data for analyses. The Explorer is the main vehicle for interactive multidimensional analysis, while Reporter adds analytical and printing capabilities. Portfolio organizes reports into presentations, and the Scheduler runs the Transformer in batch mode to update its analytical data extracts. Two of these modules, Transformer and Explorer, are worth a more in-depth discussion.

Transformer. The Transformer converts external data into a PowerPlay analytical model and one or more extract files for use in the PowerPlay Explorer. The extract file contains the data (which PowerPlay compresses), definitions for analytical dimensions such as date, region, and product line, and numeric measures such as revenue, expenses, and profit. The incoming data can be transactional records or pre-summarized totals already prepared for decision-support purposes, but the data must be a single result table. You can use Impromptu or any other query tool to produce a text file containing the input data. PowerPlay also reads Xbase files (.dbf), Paradox tables, and Lotus and Excel spreadsheet files.

After selecting an input file, the Transformer displays four multiple document interface (MDI) child windows: a dimension map across the top half of the main window, and smaller windows below that for data columns, measures, and extracts. You can rename columns, or change properties of any item in the lower windows. It's a good idea to drag and drop columns into the measures window so PowerPlay will know these are not potential dimensions. You can click on a toolbar button to make PowerPlay scan the data file and automatically generate dimensions that appear across the top of the dimension map window. Another toolbar button populates the unique values into each dimension and measure. At this point, you can double-click on a measure name to see a hierarchical display of its submeasures and the values in each level. The last essential step is to generate an extract file from the model. In addition to the default full model extract, you can create additional extracts based on subset models. For example, a subset model based on region yields an extract limited to data for that region. You can launch the Explorer to view the data in the extract.

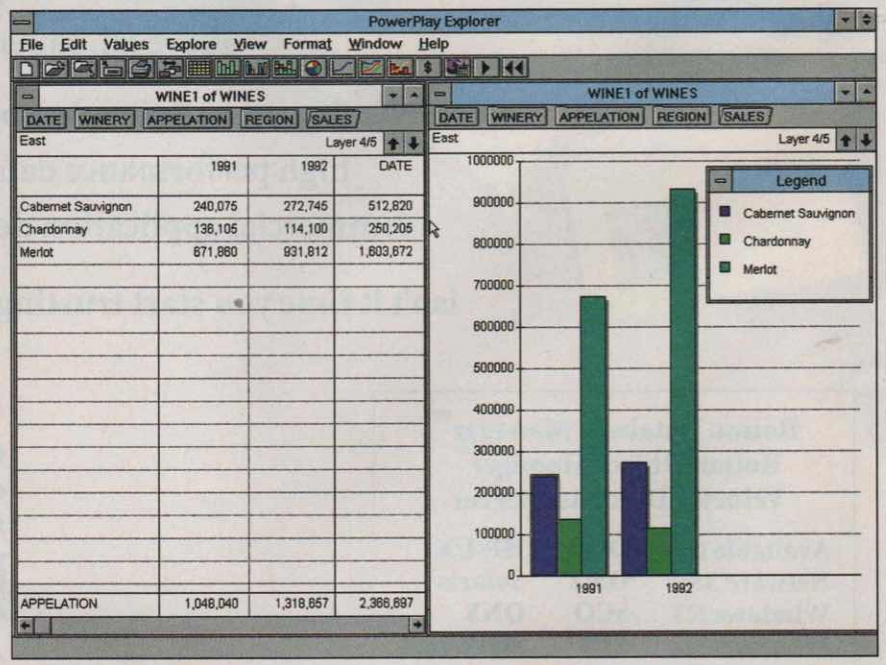
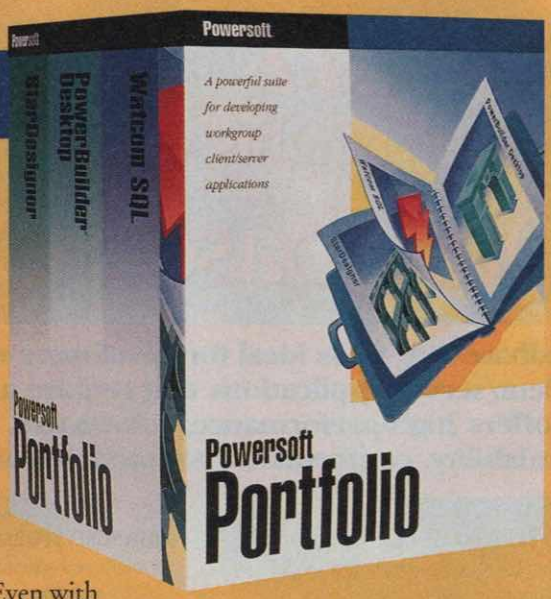


FIGURE 4 **The PowerPlay Explorer lets end users analyze data using crosstabs and eight types of graphs. The folder line across the top indicates the dimensions and measures in a PowerPlay model. Clicking on a folder displays a menu of available choices within each dimension or measure. A layer is another cross-section of a model.**

GET YOUR ASSETS IN GEAR.

Powersoft Portfolio. Everything you need to build workgroup client/server applications.



Powersoft Portfolio has everything you need to get your career in high gear. Not to mention define, build and deploy serious workgroup applications. StarDesignor™, PowerBuilder® Desktop and Watcom™ SQL 3-User were all designed from the start for client/server. And each tool is clearly recognized as the best in its class. Together they deliver the most potent client/server solutions available. At about half the retail price. With StarDesignor, data modeling is easy. StarDesignor allows you to construct new databases for high-performance applications and generate a blueprint for presentation and maintenance. PowerBuilder Desktop allows for powerful, intuitive development. Even with true object orientation and robust data access, PowerBuilder Desktop remains simple to use. Watcom SQL provides SQL database power right out of the box. It's a powerful, multi-user server that expands your development power with advanced features like stored procedures, triggers and multiple database support. So, get in gear. Get Powersoft Portfolio. To learn more about Powersoft Portfolio, call 1-800-395-3525.

$495

Powersoft Portfolio

Windows 95 Compatible

1-800-395-3525

Powersoft.

Building on the power of people.

POWERSOFT IS THE OPEN TOOLS DIVISION OF SYBASE, INC. Powersoft Corporation, (508)287-1500. http://www.powersoft.com/
Powersoft Europe Ltd., United Kingdom, +44 1494 555555. American Powersoft Corporation Pte Ltd., Singapore, +65 338-0018. Powersoft Hong Kong Ltd., Hong Kong, +852 2839 9399. Powersoft International Corporation, Australia, +61 2 9937 9600. All trademarks and registered trademarks are property of their respective owners.

CIRCLE READER SERVICE NUMBER 227

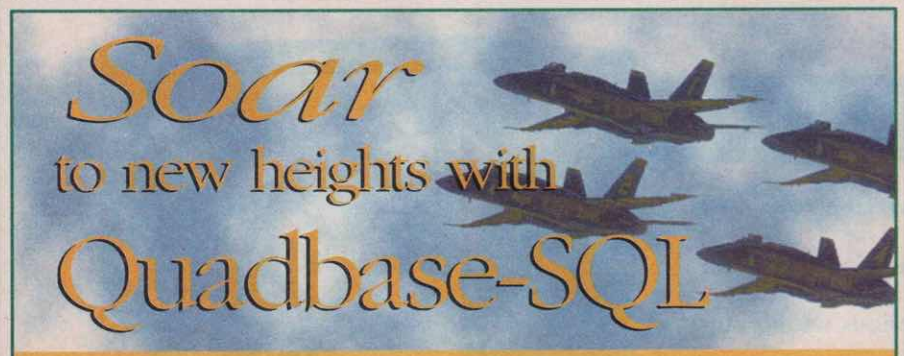
The Transformer does a pretty good — but not perfect — job of guessing which columns should be dimensions (or nested dimensions) and which should be measures. It's easy to create new dimensions and new measures (including calculated measures) and to delete or reorganize generated dimensions. PowerPlay has built-in knowledge of dates, so it generates dimensions for years, months, and quarters automatically. It also creates several relative time dimensions automatically, including current month, last month, quarter-to-date, prior quarter-to-date, year-to-date, and prior year-to-date. You can also create other special time categories, including custom fiscal years. Flexible time categories are critical to business analysts. Compared to relational databases and SQL, PowerPlay, like many other OLAP tools, offers far superior support for flexible time periods and automatic time consolidations.

In addition to using the Transformer to access data from relational databases, PowerPlay is one of the first third-party client tools that supports Arbor Software's Essbase, a multidimensional server. After establishing a connection from within the Explorer, users can double-click on dimension names and values to begin drilling down through data stored on an Essbase server.

Explorer. After you have packaged your raw data in the Transformer, the Explorer lets you view and manipulate it in several ways. (See Figure 4, page 34.) Explorer lets you open one or more views (from the same or different models), each in its own MDI child window. Each view window has a toolbar-like area across the top, but instead of buttons, you see folder icons for each dimension and measure. Below the dimension line is the data area. The initial display is a crosstab grid, complete with total lines for rows and columns. The data reflects the current dimensions you selected from the folder toolbar. Double-clicking on a dimension folder opens a filter window that lets you change the current value for any dimension. Each folder displays the name of the value you selected in that dimension.

You can also modify the crosstab by dragging and dropping new dimensions into the row or column axis, and you can swap the rows and columns using the mouse. However, I quickly realized that PowerPlay's crosstabs are pretty flat: They show only one level at a time on the rows and columns, and one measure at a time in each row and column cell. In contrast, Microsoft Excel Pivot Tables let you create complex crosstabs with multiple dimensions on each axis and multiple measures at the intersection of each row and column. Mark Morton, PowerPlay's product manager, told me that this design is intentional because it simplifies the interface. I suspect that future versions of PowerPlay will have to find a way to gain additional flexibility while remaining easy to use.

In the meantime, all is not lost. PowerPlay lets you create layers by dragging a dimension into the upper left corner of the crosstab grid. Layers act like a stack of sheets, and you navigate them using the up and down arrows above the grid. I discovered that you can open multiple copies of a view and tile the windows to see more than one layer (or display view) at a time. Speaking of display views, the crosstab grid is the initial view, but you can use toolbar buttons to see the data in one of several chart types: a simple bar graph, a clustered bar graph, a stacked bar graph, a pie chart, a single line, multiple line, or correlation graph (bars and lines).



Quadbase-SQL 4.0 is ideal for developers who build open client/server applications that require a DBMS which offers high-performance, robustness, ease-of-use, scalability, conformance to standards and low price.

HIGH-PERFORMANCE
- high-throughput transaction model - supports group commit
- asynchronous I/O
- advanced query optimizer - supports clustered index
- SMP-ready out of the box
- efficiently manages multi-gigabytes of data

HIGHLY SCALABLE
- small footprint, runs comfortably on desktop and laptop machines
- offers true client/server capabilities in workgroup environments
- capable of supporting departmental applications requiring fast response time

ADVANCED STORAGE MANAGEMENT
- file-based schemas are easy to set up and maintain
- device-based schema provide greater control over physical storage management
- on-line backup
- media recovery

EASE-OF-USE
- easy to install
- self tuned, DBA not needed for workgroup applications

CONFORMS TO STANDARDS
- entry level ANSI SQL '92
- an ODBC level 2 driver, serving as a native API to the Quadbase engine, is included

Prices start at $100. But don't be fooled by the low prices of Quadbase products. Quadbase-SQL offers many advanced features that are found (and not found) in high priced products, e.g. bidirectional scroll cursors, record locking, multiple BLOB fields per table, declarative referential integrity with full integrity options (including CASCADE on UPDATE and DELETE) and SQL 2 compliant multi-table outer joins. Platforms currently supported are Windows, Windows 95, Windows NT and NetWare. Quadbase-SQL is the DBMS that you can count on to build your most demanding 32-bit as well as 16-bit applications.

Quadbase-SQL is a trademark of Quadbase Systems Inc. All other trademarks mentioned above are the property of their respective owners.

Call for a FREE demo disk now or visit our Web site at www.quadbase.com.

Quadbase Systems, Inc.
2855 Kifer Road, #203, Santa Clara, CA 95051
Tel: 408-982-0835 • **Fax:** 408-982-0838
E-mail: sales@quadbase.com

CIRCLE READER SERVICE NUMBER 159

Armed For Analysis

As an analytical tool, PowerPlay's support for drill-down (as well as up-and-across) analysis is one of its defining features. Double-clicking on a row or column drills down to show the next level of detail. The new data replaces the old within the crosstab grid, but the total line shows the number from which you started. Double-clicking on that value rolls up to the higher-level data. Because only one level at a time is visible, you could not see and compare details for more than one higher-level group, such as years. You can also drill down and up in a graph by clicking on either the graphic element, such as a slice in a pie chart, or the associated text in the legend window.

PowerPlay provides several other analytical features. In the crosstab grid view, you can switch from viewing numeric values in each cell to percent of row or column, or as a percent of the layer or a percent of the grand totals for the view. You can highlight exceptional data by defining ranges along with a color for each criterion. You can suppress data from view (and exclude it from calculations) by formatting it as hidden, by resizing its row or column to zero, or by using the "Suppress Zero Values" option. You can impose filters by selecting a category from a floating list box (you do not need to create or enter expressions).

Other Tools

Like its Explorer sibling, the PowerPlay Reporter focuses on crosstabs and graphs rather than on lists of transactions. If you launch the Reporter directly from the Windows Program Manager, you can begin a new report. Alternatively, clicking on the Reporter button in Explorer launches Reporter and passes to it the data viewed in Explorer. The Reporter extends PowerPlay's analytical prowess by letting users define new calculations, and by creating columns that display a ranking number based on the values in a data column.

PowerPlay's Portfolio creates presentations consisting of slides onto which users place views or reports. The slides can be displayed in sequence. Slides can also have pushbuttons that jump to related slides. Because the extracts created by the Transformer are static snapshots of transactional databases, the Scheduler is an essential component because it updates the extracts and recalculates derived values so that views and reports are refreshed with current data.

Overall, I am impressed with PowerPlay's combination of ease of use and analytical depth. However, I'd like to see even more flexible display options such as multilevel hierarchical formatting of summary and detail values on the same screen.

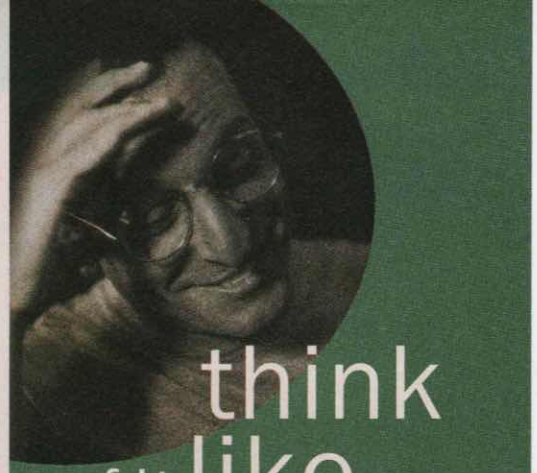
Impromptu

While PowerPlay emphasizes crosstab-style analysis, Impromptu is a report writer and query tool intended for tabular data. Impromptu's two core building blocks are data catalogs and report templates. Catalogs are prepackaged views prepared by administrators using the Administrator version of Impromptu. The User version can create reports only; it cannot create or alter catalogs.

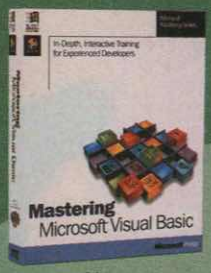
Unlike many other report writers that can connect to and access SQL databases directly, all Impromptu reports must be based on catalogs. Administrators create catalogs by connecting to a database and choosing the tables and columns that users will see. Data items are packaged into folders, and a folder can include columns from multiple tables because the administrator defines the joins, about which users never need to worry. Of course, administrators can rename columns using more business-oriented terms. Catalogs can also include filters and calculated data items. Impromptu also lets administrators print a report documenting the structure and contents of a catalog.

Impromptu accesses several client/server and desktop databases using ODBC, native drivers, and gateways. It also supports a special data format Cognos calls a "Hotfile." A Hotfile is a snapshot of a table that you can incorporate into a catalog. Hotfiles enable Impromptu to mix data from multiple databases, although the Hotfile is static, not live data. They also let users work on a report without connecting to a data source, similar to PowerPlay's extract files. You create Hotfiles by saving a report's data as a Hotfile.

A personal catalog intended for one user is the simplest type of catalog. An administrator can select a checkbox to upgrade a personal catalog into a distributed catalog. To fine-tune access by multiple users, administrators can also create classes of users, and then define filters limiting each user to a subset of the data (for example, by department or region). You can delegate administrative tasks to departmental users by changing selections in a security configuration dialog. Finally, Impromptu lets administrators activate several kinds of query governors in a catalog. You can limit the number of rows retrieved, the length of time a query can execute, allow, or prevent sorting on unindexed columns, and much more. Many end-user query tools require administrative setup, but few offer as much control as Impromptu.



think of it like a personal trainer for your code.





Now you can get the most out of Microsoft® Visual Basic® 4.0 and quickly develop applications using all of its powerful new features. All it takes is the Mastering Visual Basic CD-ROM, a self-paced tool that provides in-depth interactive training to all that's new in version 4.0.

Use its sample-code library to try out new concepts as you learn them. Then paste useful sample-code examples into your own applications. Check your knowledge with self-tests as you go. It's simply the fastest way to fully exploit the new muscle in Visual Basic 4.0.

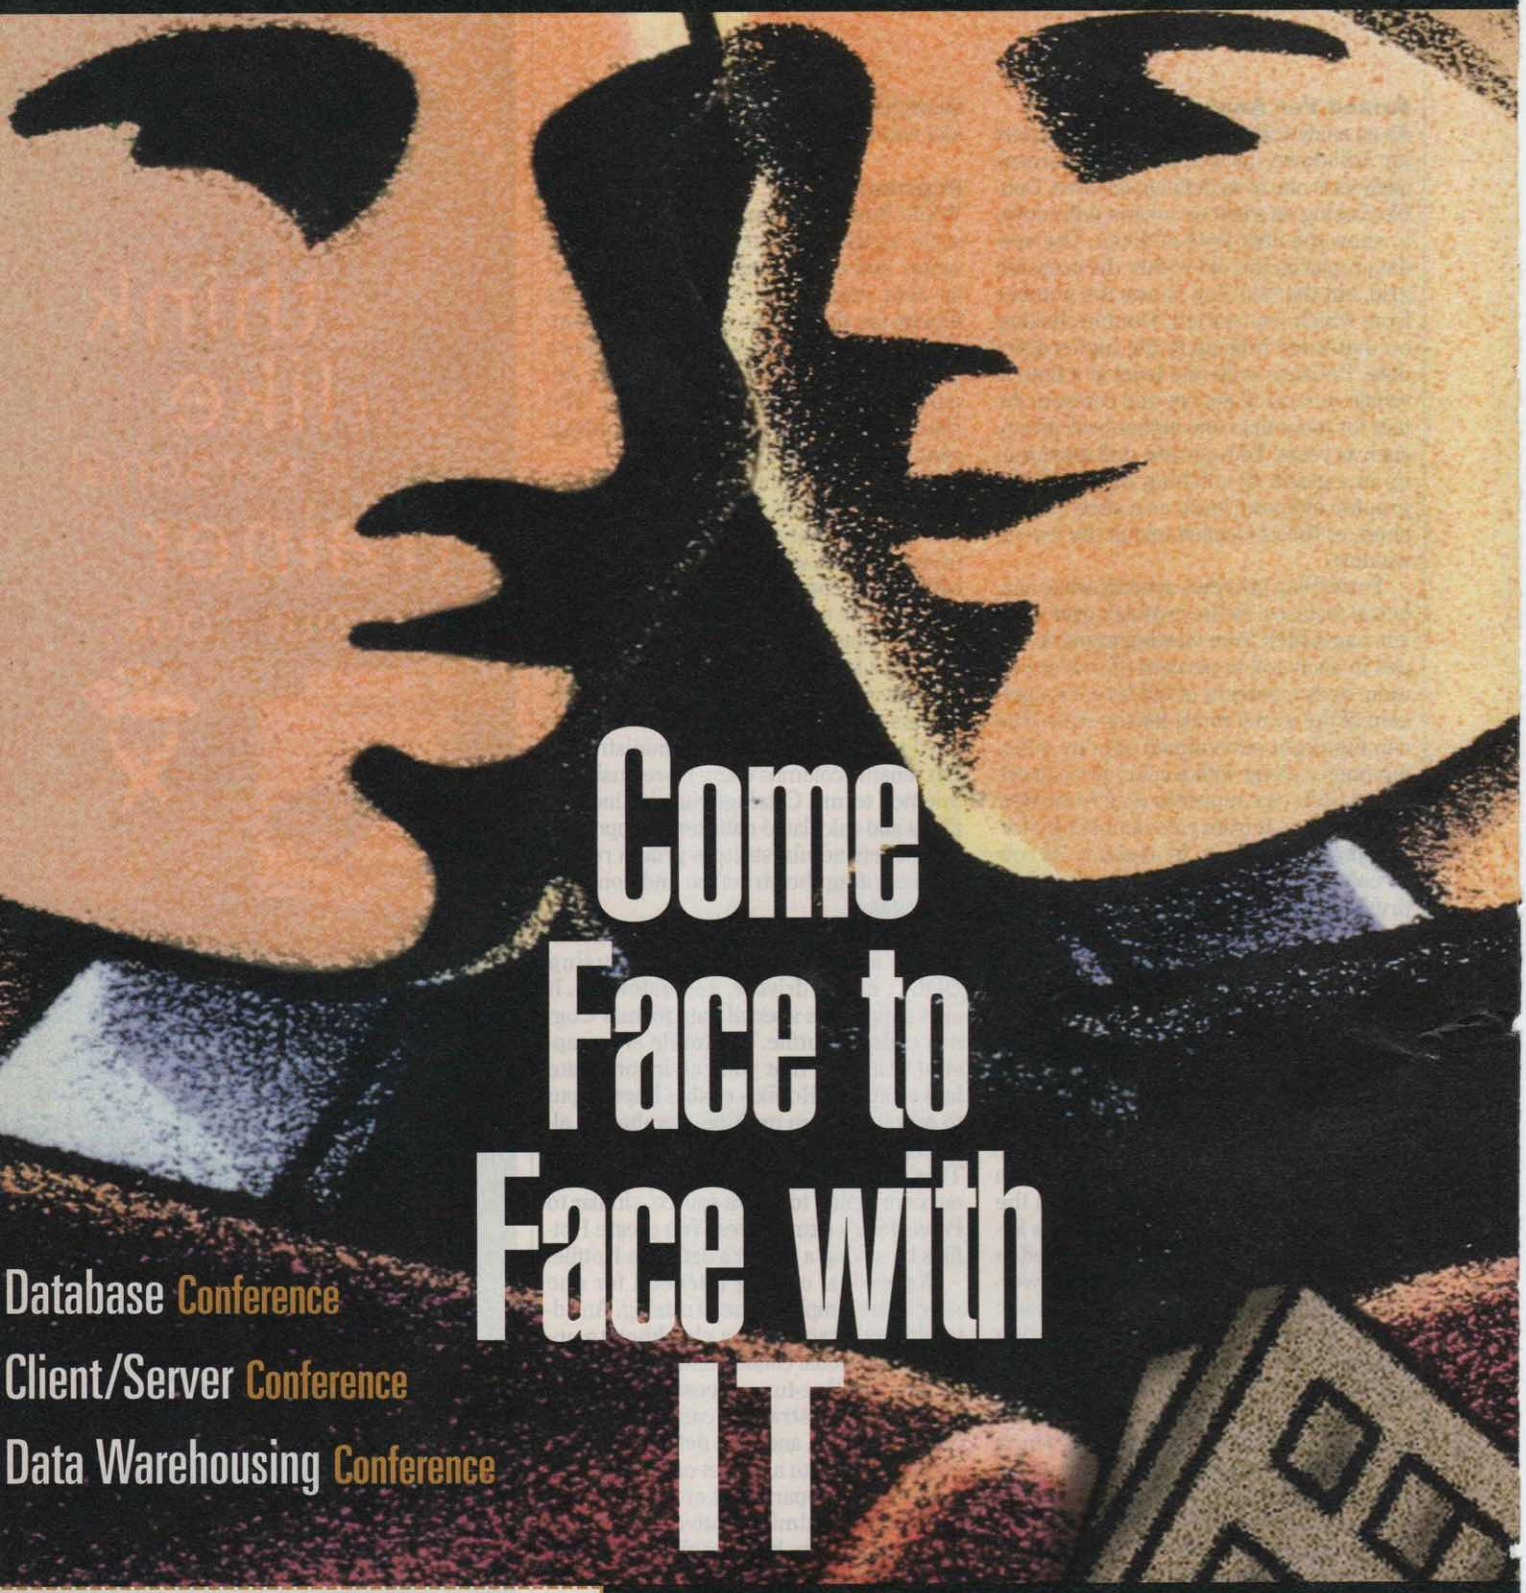
To order the Mastering Visual Basic CD-ROM, call **(800) 228-6139, Dept. XU3,** for the name of a book or software store near you or to order by phone.

Microsoft

WHERE DO YOU WANT TO GO TODAY?™

©1995 Microsoft Corporation. All rights reserved. Microsoft and Visual Basic are registered trademarks and *Where do you want to go today?* is a trademark of Microsoft Corporation.

DB/EXPO '96 SAN FRANCISCO



Come Face to Face with IT

Database Conference
Client/Server Conference
Data Warehousing Conference

Yes! I Want To Come Face To Face With The Latest In IT.

Please send me:
- ☐ My DB/EXPO '96 SF Conference Brochure
- ☐ Information on becoming an exhibitor
- ☐ FREE Expo registration information

1011SF

Name: _____ Title: _____

Company: _____

Street: _____ Mail Stop: _____

City: _____ State: ____ Zip: _____

Telephone: _____ Fax: _____

E-Mail Address: _____

BLENHEIM

24 HR FAX FORM (415) 966-8934

DB/EXPO 96

Exposition April 16-18,1996

Conference April 15-19, 1996

Moscone Convention Center

The World's Leading Database, Client/Server and Data Warehousing Exposition and Conference

FOR MORE INFORMATION CALL 1-800-2DB-EXPO

Templates

End users can create new Impromptu reports in three ways: through a ReportStyle template, through a HeadStart template, or from scratch. ReportStyle templates are nearly complete reports that contain, in addition to headers, footers, totals, and several other elements, named place holders for data items. All the user needs to do is tell Impromptu what column in a catalog should fill each place holder, and a prebuilt report is ready to run. Of course, the nature of the report has to match the user's requirements, and some customization may be necessary, but Cognos is sharp to capitalize on the fact that so many reports in so many businesses follow the same patterns. Impromptu comes with several dozen templates, and administrators can design and distribute new ones.

HeadStart templates do not contain place holders for data, but they provide place holders for report titles and corporate logo graphics. They also provide structures for common report layouts such as record listings, screen forms, combined forms and listings, and so forth. The templates already have the frame elements that users would have to create if they were building a report from scratch, but they do not have any data fields or calculations. Impromptu does not use wizards to guide users though the process of creating new reports, but it does present a single dialog with multiple tabs, and users can enter any tabbed subdialog in any order. The tab pages obtain report specifications such as the data elements drawn from a catalog, sorting, grouping, and filter conditions.

While most report writers emphasize printed output, Impromptu reports are equally at home on screen. For example, a form and list report showing customer profile information in the form, and a list of orders by that customer, can have scrollbars on the list. Impromptu also places VCR buttons in the upper-right part of the screen so that users can navigate through a report's data.

One of Impromptu's default layouts for a new report is a mailing label, but I was surprised to see that Impromptu does not provide predefined Avery label layouts. Administrators can prepare and distribute them, but that should not be necessary. Instead, they should come with the product.

Impromptu supports OLE 2, including OLE automation. The product includes CognosScript, a Basic-like macro language, but you can also use Microsoft's Visual Basic for Applications (VBA). Automation lets you run Impromptu reports from other applications and integrate Impromptu with other applications. Although powerful, this aspect of Impromptu is probably most interesting to developers and administrators rather than end users.

End-User Analysis Tools

Cognos has done a good job of balancing ease of use, robust features, and minimal administrative effort required to set up these tools for end users. Both PowerPlay 4.1 and Impromptu 3.0 are mature products that have not yet become overburdened with every imaginable feature, which is good for users who might be overwhelmed when quick results are necessary. If you need analytical and reporting tools for non-technical end users, PowerPlay and Impromptu are definitely worth careful consideration.

.....

Maurice Frank is DBMS*'s technical editor, based in Marietta, Georgia. You can reach him via CompuServe at 72167,736 or via the Internet at mfrank@mfi.com.*

We're Known for the Company We Keep

Simba Technologies–the First Name in ODBC Solutions

Simba Technologies offers a full line of SDKs: The best way to build data connectivity to your proprietary data.

Introducing SimbaExpress – Client/server access for SQL and non-SQL data.

For your ODBC needs call: **1-800-SIMBA96**

"Client-server ODBC Solutions for Business™"

SIMBA®
Simba Technologies Incorporated

2125 Western Ave.
Suite 301
Seattle, WA 98121
Tel: 206-441-0340
Fax: 206-441-9876

Formerly PageAhead Software

Apple Computer, Inc. • Attachmate • Btrieve Tec
nologies • Bull HN Information Systems • Cerebral
Technology • Cornerstone Imaging • Customer
Insight • Data Parallel Systems • Decision Suppo
Digital Equipment Corporation • ECNZ • Ergosoft
Faircom Corporation • Family Insurance Grou
Fourth Shift Corporation • Gallagher Financial
Corporation • General Electric Co. • Hewlett
Packard Company • Hyperion Software • IBM
IE Holdings BV • Ingres • Intergraph Corporati
Interleaf, Inc. • ISI • ISIS Distributed Systems
Inc. • Kvatro-Notis A.S. • Liant Software Corpor
McGraw Hill School Systems • MDBS • MDIS
Netwise • Open Data Corp. • Oracle Corporatio
Pitney Bowes • POET Software • RMS Incorporat
Red Brick Systems • Sapiens International Corp.
Samsung Electronics Co., • State of the Art, Inc.
Telekurs • Timberline Software • Unify • Unidata
Unisoft S.A. • Wall Data Inc. • WIZSOFT

SOFT NOTES

THE US OF A...CI

ACI US Inc. (Cupertino, Calif.) is shipping version 1.5 of its 4D Server client/server RDBMS product. Based on ACI's 4th Dimension database, the tool integrates client software, the server application, and network modules into one product. The platform-independent architecture lets client software and database applications run on either Microsoft Windows or Apple Macintosh operating systems without any modifications.

The 4D Server ships with a set of development tools, including 4D Compiler, which translates a 4D application into fast-executing machine code. The product also supports team development by letting multiple developers access the same database application at the same time; any changes made to the application are sent to the client workstations automatically. Supported network protocols include TCP/IP, IPX/SPX, and ADSP (AppleTalk); 4D Server can support one or all of these protocols simultaneously.

The Windows version runs on a 386, 486, or Pentium machine with Windows 3.1 or later, and at least 8MB of RAM; or Windows 95 with 8MB of RAM, or Windows NT Workstation or Server with at least 16 MB of RAM. The Macintosh version requires a Macintosh with System 6.0.7 or later, and a minimum of 2MB of RAM. 4D Server starts at $1295, and the 4D Server 1.5 Macintosh upgrade is $495. *Contact ACI US Inc., 20883 Stevens Creek Blvd., Cupertino, CA 95014; 800-881-3466, 408-252-4444, or fax 408-252-4829.*

SHORT NOTES

Sterling Software Inc. (Atlanta) released its Key for Workgroup suite, a single-vendor client/server application development suite that includes tools for workflow modeling, business process reengineering, process management, application analysis, and visual development. The suite consists of six integrated products: Key:Advise, Key:Model, Key:Empower for Visual Basic, Key:Empower for PowerBuilder, and Key:Assemble. *(Contact Sterling Software, Inc. Applications Development Division, 800-444-8575, 404-231-8575, or fax 404-364-0267.)* **BMC Software Inc.** (Houston) made several additions to its DB2 product family. New are SQL Recall, which stores certain repetitive SQL queries in cache and accesses them only as needed; Logmax for DB2, which enhances administrator access to the DB2 log; Exceptor for DB2, which identifies and corrects DB2 problems before business processes are affected; and version 3.0 of Data Packer for DB2. Also shipping is Paclog, a tool that reduces the size of DB2 logs. *(Contact BMC Software Inc., 800-841-2031, 713-918-8800, or fax 713-918-8000.)*

Good news for companies struggling with implementing an enterprise-wide systems management strategy: **Platinum Technology Inc.** (Oakbrook Terrace, Ill.) announced the Platinum Repository/Open Enterprise Edition. The client/server toolset includes data administration and enterprise-wide metadata management tools. It also serves as a universal relational model for heterogeneous RDBMSs to interact and share information, and distributes resources to the appropriate platforms via a three-tiered client/server architecture. *(Contact Platinum Technology Inc., 708-620-5000 or fax 708-691-0710.)*

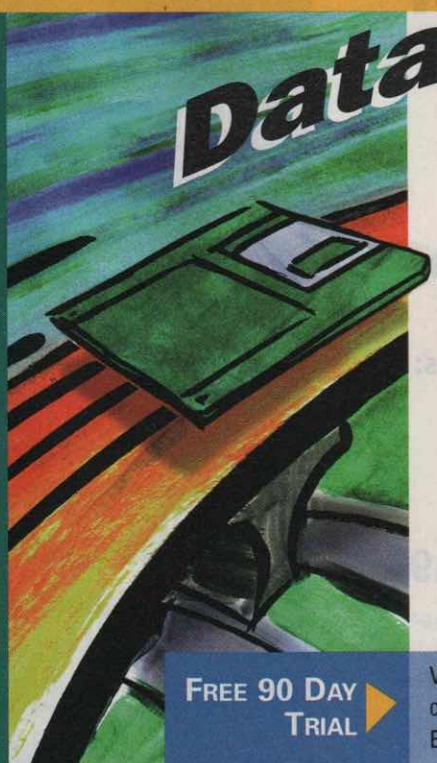
On the data warehousing front, **Thomson Software Products** (Norwalk, Conn.) announced RP/Server, an enhanced version of the company's earlier Must 1.0 prod-

RUN YOUR EXISTING ODBC DATABASE APPLICATION ON THE INTERNET

IT'S SECURE! IT'S EASY! IT'S NOW!

DataRamp

- RELATIONAL: CONNECTS TO THE MOST POPULAR RELATIONAL AND PERSONAL DATABASE SYSTEMS.
- ODBC: WORKS WITH ANY WINDOWS APPLICATION THAT ACCESSES DATABASES.
- EASY TO INSTALL: CLIENT AND SERVER CAN BE INSTALLED IN 5 MINUTES.
- SERVER MONITOR: SHOWS ACTIVITY ON THE SERVER.
- PLATFORMS: DATARAMP CLIENT ON WINDOWS, WINDOWS95, WINDOWS NT. DATARAMP SERVER ON WINDOWS NT.
- SECURE: USES RSA'S PATENTED RC4 ENCRYPTION ALGORITHM FOR ALL NET COMMUNICATION. RSA



DATARAMP SERVER
DATARAMP CLIENT
THE INTERNET

- VISUAL BASIC
- POWERBUILDER
- VISUAL C++
- MICROSOFT ACCESS
- CRYSTAL REPORTS
- CA-OPENROAD
- ORACLE
- SQL SERVER
- CA-INGRES
- dBASE
- INFORMIX

FREE 90 DAY TRIAL

VISIT OUR WEB SITE: http://DataRamp.com
CALL (617) 863-2339 OR FAX (617) 863-6220
EMAIL: sales@DataRamp.com

DataRamp
From Working Set, Inc.

© Working Set, Inc.
All product names trademarked by their respected companies.

CIRCLE READER SERVICE NUMBER 121

SQA Suite™ Seamlessly-Integrated Testing For Enterprise Client/Server

SQA Suite™

SQA Process SQA Robot SQA Manager SQA LoadTest

Highlights

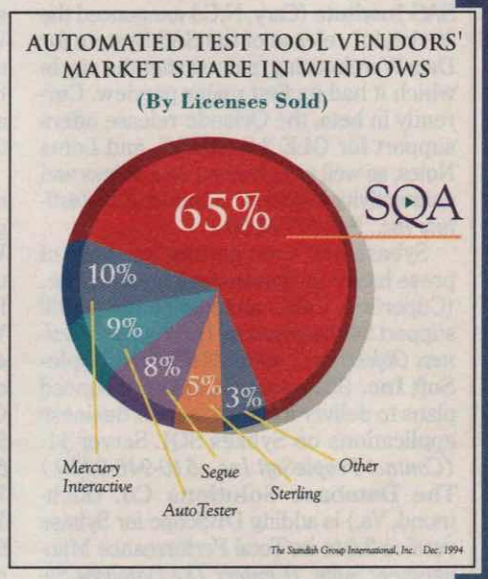
- Comprises the industry's only seamlessly-integrated suite of test tools for **comprehensive, end-to-end testing of enterprise-level Windows client/server applications**
- Uses Object Testing™ to test client/server applications developed with PowerBuilder, Visual Basic, SQLWindows, and other development environments for **complete, object-level testing of client/server applications**
- Combines the speed of Object-Oriented Recording™ with an integrated, extensible programming environment for **automatic and rapid test creation with the full power of programming**
- Offers load, stress, and multi-user testing of client/server applications on TCP/IP, NETBIOS, or IPX/SPX networks to **ensure Windows client/server applications are production-ready before deployment**
- Provides an integrated, open network test repository and advanced test management and measurement tools for **quantifying results, managing defects, and coordinating work among members of the project team**

SQA®

THE LEADER IN WINDOWS CLIENT/SERVER TESTING

The Leading Windows Client/Server Test Environment

SQA Suite™ is the only seamlessly-integrated product suite for the automated testing of enterprise-level Windows™ client/server applications. Built on an integrated, open network test repository, *SQA Suite* combines leading-edge client/server testing power, comprehensive management tools, and a formal methodology to set the standard for automated testing of Windows client/server applications. And with Object-Testing™, you'll go beyond testing your application's GUI with the most advanced object-level integration available for PowerBuilder, Visual Basic, and other client/server development tools. That's why SQA is the leader in the testing of Windows client/server applications, and SQA's products have won more awards than any other test tool on the market.



"SQA is the #1 vendor of test tools for Windows, the strategic client market in mission-critical client/server applications."

Jim Johnson, Chairman,
The Standish Group International, Inc.

"...well integrated suite..." "...the hands-down winner..." "...a blockbuster..." "...indispensable..." Best Client/Server Test Tool

For a FREE SQA Demo Disk or Seminar Invite call 800.228.9922

Email us at info@sqa.com or seminars@sqa.com

SQA, Inc. 10 State Street Woburn, MA 01801 Tel 800.228.9922 or 617.939.3000 Fax 617.932.3280

CIRCLE READER SERVICE NUMBER 76

uct. RP/Server is a middleware tool that offers mainframe-to-desktop data delivery for companies building a data warehouse or migrating to a client/server environment. *(Contact Thomson Software Products, 203-845-5000 or fax 203-845-5252.)* **Information Builders** (New York) announced a Unix-based version of its Enterprise Copy Manager, which automates the entire process of building and maintaining a data warehouse, and offers data extraction, transformation, and transportation facilities. *(Contact Information Builders Inc., 212-736-5533 or fax 212-268-7470.)* Finally, **SAS Institute** (Cary, N.C.) announced the "Orlando" release of the SAS System for Data Warehousing, named after the city in which it had its first major preview. Currently in beta, the Orlando release offers support for OLE 2.0, ODBC, and Lotus Notes, as well as enhanced data access and connectivity facilities. *(Contact SAS Institute Inc., 919-677-8000.)*

Sybase has been getting its share of press lately as **Business Objects Inc.** (Cupertino, Calif.) announced that it will support Sybase System 11. *(Contact Business Objects Inc., 408-973-9300.)* **PeopleSoft Inc.** (Pleasanton, Calif.) announced plans to deliver its client/server business applications on Sybaes SQL Server 11. *(Contact PeopleSoft Inc., 510-946-9460.)* **The Database Solutions Co.** (Richmond, Va.) is adding DBScope for Sybase version 2.0 to its Total Performance Management Suite. *(Contact The Database Solutions Co., 800-933-7668, 804-794-0354, or fax 804-794-0813.)*

SelectStar is shipping version 1.5 of its relational OLAP tool StarTrieve. Oriented toward mainstream business users, the product features enhanced drill-down, data ranging, and data grouping capabilities. *(Contact SelectStar, 212-243-8413 or fax 212-627-9784.)* Another OLAP tool provider, **Kenan Technologies Inc.** (Cambridge, Mass.), announced the availability of Acumate ES version 1.3. New are an upgraded multidimensional database, the Acutrieve analysis and reporting application, and support for new server platforms including IBM's RS/6000 running on AIX. Acumate ES clients now support Windows 95. *(Contact Kenan Technologies Inc., 617-225-2224 or fax 617-225-2291.)* For those looking for more information on OLAP itself, **Business Intelligence Inc.** (Norwalk, Conn.) has released *The OLAP Report*, an independent and comprehensive guide to OLAP applications, technologies, and products that are offered in the United States and Europe. *(Contact Business Intelligence Inc., 800-690-6527, 203-899-7920, or fax 203-838-4350.)*

IN PRINT



WORKING WITH WINDOWS

Frequent *DBMS* contributor, consultant, database programmer, and author Ken North has released a new book that offers help for developers who create and use Windows-based database applications. *Windows Multi-DBMS Programming: Using C++, Visual Basic, ODBC, OLE 2, and Tools for DBMS Projects* offers a comprehensive survey of DBMS engines and servers, Windows software development products, multidatabase programming, database applications, ODBC, data access and presentation tools, and client/server tool suites.

The book ships with a CD-ROM or disk of Windows-based multi-DBMS programming companion software including Windows Q+E Database Library and MultiLink/VB, Quadbase-SQL 3.0a, Crystal Reports for Windows, and Microsoft Video for Windows Runtime System. *Windows Multi-DBMS Programming* (book/disk set ISBN 0-471-01675-6; book/CD-ROM set ISBN 0-471-01676-4) costs $49.95 for the book and disk, and $59.95 for the book and CD-ROM. *Contact John Wiley & Sons Inc., 605 Third Ave., New York, NY 10158-0012; 800-225-5945, fax 908-302-2300, email compbks@jwiley.com, or via the World Wide Web at http://www.wiley.com.*

HELP FOR MANAGERS

As client/server becomes more of a fact of life in businesses, many IS managers are struggling to understand the beast that has taken over their organizations and doesn't seem to be going away. A good reference book for such individuals is Karen Watterson's *Client/Server Technology for Managers*. Short and to the point, the book begins with a clear explanation of what client/server is and why it is so important, and then dissects client/server to study the client, server, network, middleware, object, open system, and standards components.

Watterson gives advice on how to determine the cost of client/server (with troubleshooting tips and recommendations from people out in the field), and looks at where information technology is going and how you can make decisions that will withstand future changes. *Client/Server Technology for Managers* (ISBN 0-201-40920-8) costs $17.95. *Contact Addison-Wesley Publishing Company, One Jacob Way, Reading, MA 01867; 617-944-3700 or fax 617-942-1117.* ■

EVENT ALERTER

JANUARY

- **Lotusphere '96,** Orlando, Fla., January 21–25. **Contact** Lotusphere '96, 800-655-6887.
- **USENIX 1996 Technical Conference,** San Diego, January 22–26. **Contact** USENIX Conference Office, 714-588-8649 or fax 714-588-9706.
- **ComNet '96,** Washington, January 29–February 1. **Contact** Mitch Hall Associates, 617-361-8000 or fax 617-361-9074.
- **SoftExpo '96,** Santa Clara, Calif., January 30–February 1. **Contact** Webcom Communications, 303-745-5711.

FEBRUARY

- **Data Warehousing Conference,** Orlando, Fla., February 6–8. **Contact** DCI, 508-470-3880 or fax 508-470-0526.
- **UniForum '96/Enterprise Computing Solutions,** San Francisco, February 12–16. **Contact** Softbank Comdex Inc., 617-433-1500 or fax 617-449-2674.
- **Networks Expo Boston,** Boston, February 13–15. **Contact** Blenheim Group USA, 800-829-3976, 201-346-1400, fax 201-346-1532, or http://www.shownet.com.
- **Mactivity,** Boston, February 13–15. **Contact** Mactivity Inc., 408-354-2500 or http://www.mactivity.com.
- **Client/Server Project Management,** Chicago, February 26-27. **Contact** DCI, 508-470-3880 or fax 508-470-0526.
- **Decorum '96,** Tarpon Springs, Fla., February 27–29. **Contact** Transarc Corp., 412-338-6900; e-mail: decorum@transarc.com; or http://www.transarc.com.

MARCH

- **Smalltalk Solutions,** New York, March 4–7. **Contact** SIGS Conferences, 212-242-7447, fax 212-242-7574, e-mail: info@sigs.com, or http://www.sigs.com.
- **Software Development '96,** San Francisco, March 25–29. **Contact** Miller Freeman Inc., 415-905-2210.
- **Database and Client/Server World,** Boston, March 26–28. **Contact** DCI, 508-470-3880 or fax 508-470-0526.
- **TopSpeed Developer's Conference '96,** Orlando, Fla., March 26–28. **Contact** TopSpeed Developer's Conference, 800-767-2336 or fax 508-470-0526.

where would the most powerful development tool be without the recycling genius of edwin hoogerbeets?

If Edwin Hoogerbeets were a developer, he'd be a proud man. Because new Microsoft® Visual C++ 4.0™ is based on the very principle he lives by: Never create from scratch.

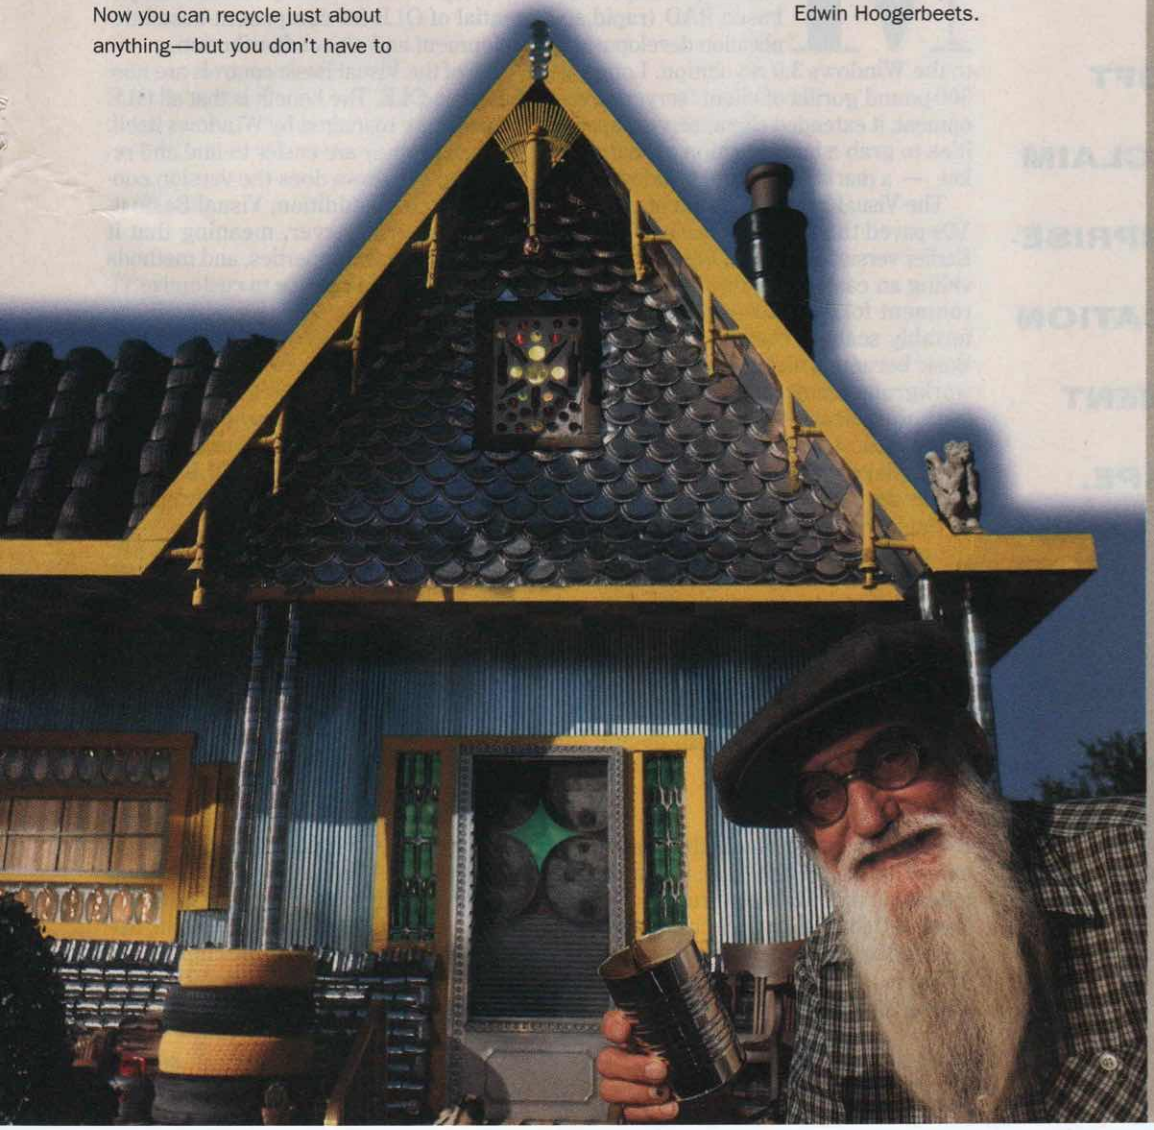
Who is Edwin Hoogerbeets? A recycling genius who sees every used object as an opportunity for creative brilliance. The same way a developer might if he or she used new Visual C++ 4.0. No other development tool offers such reusability, allowing you to deliver more robust applications in less time.

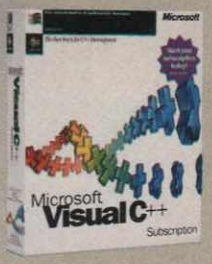
Now you can recycle just about anything—but you don't have to be a genius to do it. The new Component Gallery allows you to store and reuse your own C++ components and OLE controls, as well as 3rd-party components. And, MFC 4.0 gives you more than 150 classes and 120,000 lines of code you don't have to write or test.

The new Developer Studio makes reuse intuitive. With ClassView, you can view the relationships between classes. You'll even have one-click access to the MSDN Development Library, Microsoft Visual Test, Fortran PowerStation, and Microsoft Visual SourceSafe™ version control system. And of course, expanded C++ language support, including namespaces and RTTI, gives you added flexibility.

Now that you have a better idea of what's possible when you recycle, start with a Visual C++ Subscription, which includes the Visual C++ development system 4.0, plus three additional Subscription Releases shipped automatically.

For more information about the Visual C++ 4.0 Subscription, or to get your free Developer Roadmap CD, call (800) 719-5577, Dept. 67T. You can also get information at **http://www.microsoft.com/visualc** or by calling one of the resellers listed below. Who knows? You just might become the next Edwin Hoogerbeets.





Component Gallery. Quick access to OLE controls and reusable C++ components.

MFC 4.0 Library. Adds latest support for Windows® 95 controls.

Client-Server Support. Built-in Jet database engine and ODBC support to access enterprise data.

Custom AppWizard. Now create your own AppWizard for your unique application needs.

Multiple Platform Support. Target Intel,® RISC,* and Macintosh®* platforms with a single codebase.

ASAP Software Express
(800) 883-2911

Programmer's Shop
(800) 446-1185

Programmer's Paradise
(800) 892-1187

Software Spectrum
(800) 840-8194

Microsoft
WHERE DO YOU WANT TO GO TODAY?®

*Versions of Visual C++ 4.0 for RISC and Macintosh available separately. ©1995 Microsoft Corporation. All rights reserved. Microsoft and Windows are registered trademarks and Visual C++, Visual SourceSafe and *Where do you want to go today?* are trademarks of Microsoft Corporation. Intel is a registered trademark of Intel Corporation. Macintosh is a registered trademark of Apple Computer, Inc.

VISUAL B

By David S. Linthicum

MICROSOFT STAKES ITS CLAIM IN THE ENTERPRISE-WIDE APPLICATION DEVELOPMENT LANDSCAPE.

Microsoft Visual Basic was one of the first tools that brought drag-and-drop component-based RAD (rapid application development) to the Windows 3.0 revolution. Long the 300-pound gorilla of client/server development, it extended client/server capabilities to grab a large portion of that market — a market it still holds today.

The Visual Basic revolution of the early '90s paved the way for today's evolution. Earlier versions of Visual Basic, while providing an easy-to-use development environment for Windows, could not comfortably scale to the enterprise. Visual Basic became a niche player in RAD and workgroup client/server development. With the release of its latest rendering, Visual Basic 4.0, Microsoft may prove that it is a player in the multi-tiered client/server development market.

But is Visual Basic 4.0 really ready for the enterprise? And, more important, is the enterprise ready for Visual Basic 4.0? Let's explore the new features of Visual Basic 4.0, including its new reliance on OLE, latest database features, language enhancements, and new application partitioning capabilities for multi-tiered development; then you can decide for yourself.

OLE Evolution

Of the many new features that you would expect with a new release, the most significant is that Microsoft rewired Visual Basic 4.0 for OLE (Object Linking and Embedding). For the last two years OLE, which is a standard of "interoperation" among applications and application components, has been the mantra of Microsoft. It was not until the release of Windows 95, and now Visual Basic 4.0, that developers could finally release the potential of OLE for component-based development and object distribution.

All of the Visual Basic controls are now based on OLE. The benefit is that all OLE controls are managed by Windows itself. Therefore, they are easier to find and reuse, and Windows does the version control for you. In addition, Visual Basic itself is an OLE server, meaning that it exposes objects, properties, and methods that applications can use to customize Visual Basic.

With Visual Basic 4.0, you can also write OLE servers that can provide functionality to other applications. Later in this article, I describe how to create OLE servers. Having created part or all of a Visual Basic application as an OLE server, other developers can use it locally or over a network.

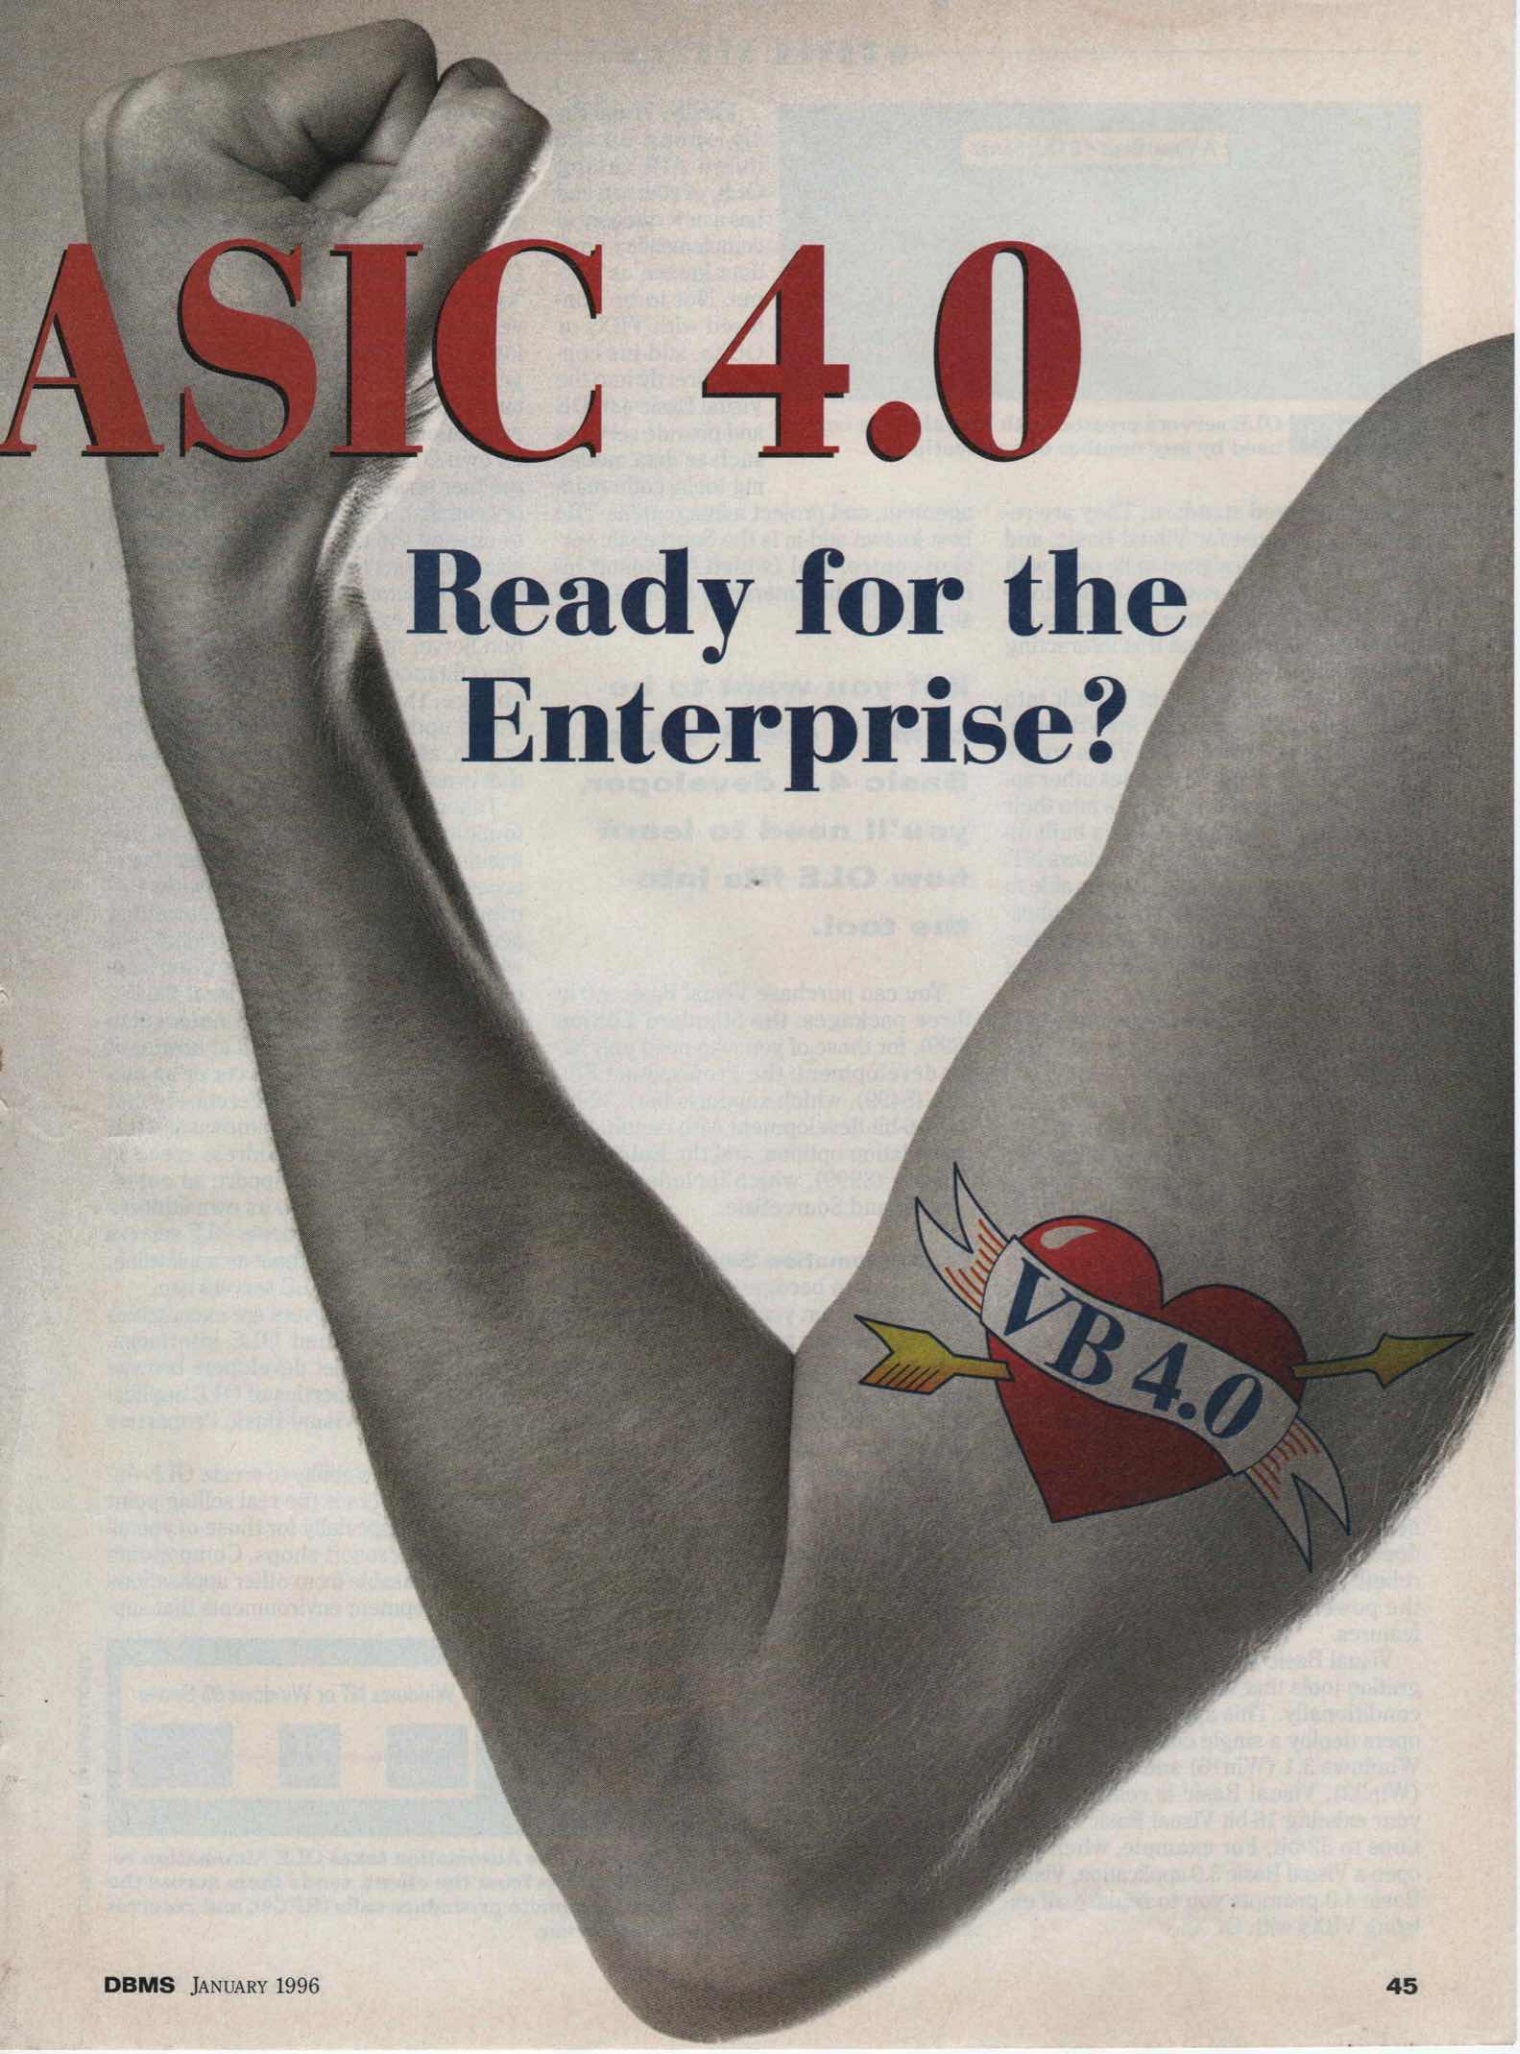
Developers can now create and distribute OLE objects and share them with other applications on the network. Through these distributed OLE objects (using a variety of object-access and data-access tools), Visual Basic 4.0 finally supports multi-tier client/server development.

The widely popular VBXs (Visual Basic controls) are no longer the staple of Visual Basic. OCXs (OLE custom controls), which use subsets of OLE technology, replace VBX building blocks for applications. Although VBXs have been adopted in other development environments such as PowerBuilder, Delphi, Microsoft Visual C++, and Borland C++, many developers consider the VBX standard a kludge. (VBXs are not based on a

David S. Linthicum is a technical manager with EDS in Falls Church, Virginia. He's the author of several books on software development, a speaker, and an associate professor of computer science at a local college. You can reach David via CompuServe at 70742,3165.

ASIC 4.0

Ready for the Enterprise?



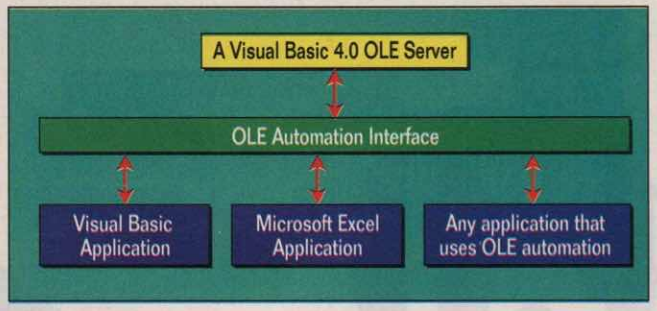


FIGURE 1 **OLE servers created with Visual Basic can be used by any number of applications.**

widely accepted standard. They are really designed just for Visual Basic, and they were never designed to be used with other development environments. Moreover, every VBX that interacts with its environment assumes that it is interacting with Visual Basic.)

The OLE standard that is built into OCXs is much broader than the VBX standard (OCXs are 32-bit, while VBXs are 16-bit). The OLE standard enables other applications to incorporate OCXs into their environments, because OLE is built directly into Windows 95 and Windows NT. Plus, the operating system itself is able to track OLE components for the applications implementing them, whereas the operating system does not keep track of VBXs.

Microsoft has made a firm commitment to an OLE-based future. As a result, hundreds of companies that currently sell VBXs are now scrambling to rewrite their VBXs as OCXs for the new market. Fortunately, Microsoft provides a migration toolkit that makes the move relatively painless. If you've written your own VBXs, it's time to think about converting.

But OLE is only part of the story. With release 4.0, developers can create DLLs entirely within Visual Basic, thereby avoiding clumsy 3GLs. There are several database improvements as well, from replication and "fat" cursors to enhanced referential integrity. (Using "fat" cursors, developers can customize the size of the local data cache to maintain more data from an answer set on the client, thus reducing network traffic. However, this approach does require more local resources.) The rebuilt 32-bit JET 3.0 database engine is the power behind these new database features.

Visual Basic also comes with new migration tools that let developers compile conditionally. This approach lets developers deploy a single code base on both Windows 3.1 (Win16) and Windows 95 (Win32). Visual Basic is ready to take your existing 16-bit Visual Basic applications to 32-bit. For example, when you open a Visual Basic 3.0 application, Visual Basic 4.0 prompts you to replace all existing VBXs with OCXs.

Finally, Visual Basic opens up the menu API (using OLE, of course), and has a new category of complementary products known as add-ins. Not to be confused with VBXs or OCXs, add-ins connect directly into the Visual Basic 4.0 IDE and provide services such as data modeling tools, code management, and project management. The best known add-in is the SourceSafe version-control tool (which Microsoft includes with the Enterprise addition of Visual Basic).

If you want to become a great Visual Basic 4.0 developer, you'll need to learn how OLE fits into the tool.

You can purchase Visual Basic 4.0 in three packages: the Standard Edition ($99), for those of you who need only 32-bit development; the Professional Edition ($499), which supports both 32-bit and 16-bit development with conditional compilation options; and the Enterprise Edition ($999), which includes a code profiler and SourceSafe.

OLE Automation Servers

If you want to become a great Visual Basic 4.0 developer, you'll need to learn how OLE fits into the tool. The success of Visual Basic 4.0 is directly linked to the success of OLE, even though OLE has been difficult to implement thus far. The API is difficult to use and understand, and OLE support among tools has been erratic.

But Visual Basic 4.0 could be the first development tool that makes OLE a reality for most organizations. OLE is also the middleware layer for distributed Visual Basic applications using Distributed OLE. Moreover, the OLE technology lets developers encapsulate things such as legacy code and business processes into OLE Automation Servers, which are reusable across the enterprise.

Automation Servers can be single- or multi-user, and they can provide services requested by other objects, either locally or over a network. Basically, any portion of an application that you need to share with other applications is a good candidate for OLE Automation Servers. In a nutshell, OLE Automation Servers let developers "componentize" applications. This lets developers isolate chunks of application logic, which can later be modified independently. For example, developers of a tax application may want to place the tax calculation logic that changes yearly into its own OLE Automation Server, which can then share several applications (locally or centrally). Therefore, if developers need to change the tax logic for all client applications, they would have to change only one OLE Automation Server.

Another example is an OLE Automation Server that is set up to monitor real-time financial feeds. As market prices change, the OLE Automation Server would update itself with the new information, as well as update any application that is using the Automation Server.

I should mention here that an OLE Automation Server is not the same as Distributed OLE. Rather, it is a server that is accessible locally or over a network. Distributed OLE is an OS-level feature that performs many of the same functions, but will incorporate more features when it appears in future releases of Visual Basic.

When you create an OLE Automation Server, you have the option of turning it into an OLE in-process server or an out-of-process server. The difference is that in-process servers (also known as OLE DLLs) share the same address space as the applications they support; an out-of-process server exists in its own address space. That's why in-process OLE servers can't be distributed without an application, and out-of-process OLE servers can.

Out-of-process servers are executables that have predefined OLE interfaces. These interfaces let developers browse methods and properties of OLE applications using the Visual Basic Properties browser.

Visual Basic's ability to create OLE Automation Servers is the real selling point of this tool, especially for those of you already in Microsoft shops. Components are now reusable from other applications and development environments that sup-

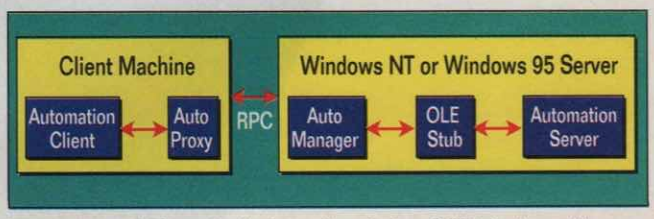


FIGURE 2 **Remote Automation takes OLE Automation requests from the client, sends them across the network using remote procedure calls (RPCs), and receives them on the server.**

ILLUSTRATIONS BY BARBARA LICATA

port OLE. (See Figure 1, page 46.) Creating OLE Automation Servers is easy. You begin by adding a new class module to the project (for object definition). Then you define the methods and properties to be encapsulated inside of the object. Next, you set the "Class Public" property to True, and define Project StartMode to OLE Server. Although it sounds complex, I created a simple OLE Automation Server in just a few minutes.

After you define your OLE Automation Server as an in-process or out-of-process server, you can test it. After chasing all the bugs in the standard modules, you begin another instance of Visual Basic. After that, you must reference the OLE Automation Server using References from the Tools menu. Then it's time to create a test control application, and test the methods and properties encapsulated inside the OLE Automation Server.

Debugging OLE components is easy with Visual Basic because of the new enhancements in error-handling. The Err object now provides developers with more information, such as source applications, descriptions, and online help.

Visual Partitioning

The real benefit of out-of-process servers is that they can execute on remote machines and are accessible over a network, thus enabling Visual Basic's new application partitioning capabilities. Remote Automation takes OLE Automation requests from the client, sends them across the network using remote procedure calls (RPCs), and receives them on the server. (See Figure 2, page 46.) The distributed Automation Server objects communicate with one another using asynchronous notification, or through OLE callbacks. Also, flexible protocol support is provided via TCP/IP, NetBEUI, IPX/SPX, and Named Pipes. Using this model, any Visual Basic application can pass an object reference to any remote OLE server that is accessible over the network. The idea is that Visual Basic 4.0, with Windows 95, will finally provide the back-end application processing that was missing from earlier Microsoft client/server products.

If you don't need to distribute OLE Automation Servers (that is, if you use a single machine), then an in-process server (or OLE DLL) will do just fine because it's addressable by a local application only. The in-process servers are automatically loaded into the same address space as the application, which results in better performance.

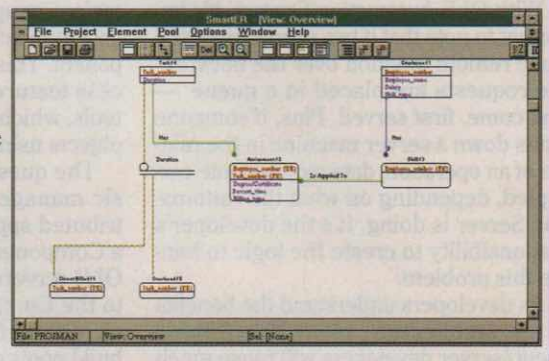
This does not mean that you can't use an in-process server on remote machines. However, a remote application must load the in-process server into its local address space to make it accessible to other applications on the network. In contrast, out-of-process servers don't need a host application. Plus, while out-of-process servers are useful for application partitioning, they are also useful for delivering applications in discrete pieces (on a single machine). By having out-of-process servers, you can deliver pieces of application code that will not interfere with other application servers; also, developers can revise them more quickly and easily.

Using this multi-tier client/server model, developers can place several business processing functions contained inside OLE Automation Servers on any available computer. Thus, updating applications is simply a matter of changing the logic in a shared distributed object. Referring back to the tax application example, you could place tax law logic inside OLE Automation Servers that are accessible over the network from several Visual Basic front-end applications. As the laws change, you can update the logic in the central OLE Automation Servers, and thus propagate the new functionality (assuming that parameter interfaces remain the same).

Make the SmartER™ choice . . .

SmartER from KBSI for all your data modeling, database design, and SQL code generation needs.

SmartER is a powerful tool for any enterprise needing to analyze its information and database structures and design new database systems. With SmartER, your team has the power to easily and effectively:



- Determine information requirements and document physical models.
- Design relational databases that can support system implementation.
- Automatically generate SQL code for database implementation.
- Reverse engineer existing databases into representative data models by importing SQL.

SmartER shares information with other KBSI tools and integrates information through its ODBC interface. This integration supports the generation and reverse engineering of database structures in most ODBC-compliant systems.

For more information, contact Knowledge Based Systems, Inc.

WWW: http://www.kbsi.com/ • E-mail: products@kbsi.com

One KBSI Place • 1500 University Drive East • College Station, TX 77840
Phone: 409.260.5274 • Fax: 409.260.1965

Crystal Plaza One • 2001 Jefferson Davis Hwy, Suite 600 • Arlington, VA 22202
Phone: 703.415.8600 • Fax: 703.415.8601

CIRCLE READER SERVICE NUMBER 84

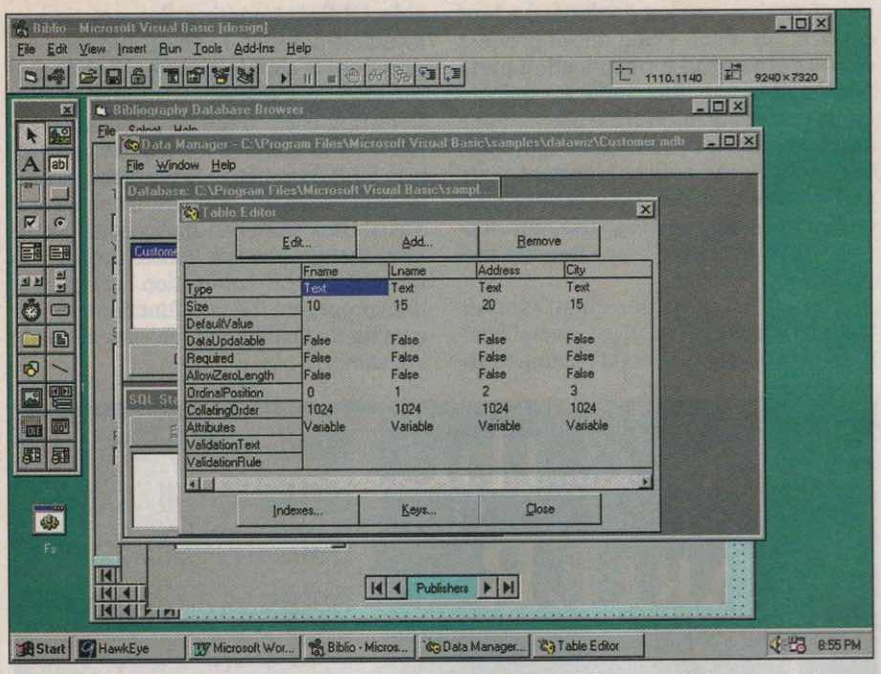


FIGURE 3 **Visual Basic's Database Manager provides complete control over any database to which Visual Basic connects.**

With OLE Automation Servers, it's important to note that if two users access the same remote function over the network, the requests are placed in a queue — first come, first served. Plus, if someone shuts down a server machine in the middle of an operation, data may become corrupted, depending on what the Automation Server is doing. It's the developer's responsibility to create the logic to handle this problem.

As developers understand the benefits of this architecture, many Visual Basic client/server developers will move much of the logic that currently resides on the database server as stored procedures and triggers to distributed Visual Basic OLE Automation Servers. How you would set this up really depends on the application. For instance, if you have a lot of users going directly to the database (such as with a DSS application), you should place your integrity constraints there. In a transaction-processing environment, you could use the middle tier (for example, OLE servers) to funnel the transactions to the database. Moreover, because this is an open architecture, any tool that uses OLE (such as Microsoft Access, Visual FoxPro, PowerBuilder, Excel, and so on) can use the Visual Basic OLE Automation Servers as well.

Visual Basic 4.0 also provides several tools for managing remote data access, including a remote data controller. This data controller is similar to the JET database engine found in version 3.0. However, the controller provided with Visual Basic 4.0 can populate data-aware controls for ODBC data sources. Using this mechanism, developers may place an interface wrapper around the ODBC API, and thus access the database as a component. This is similar to the database access features found in Smalltalk-based tools, which treat relational databases as objects using wrappers.

The question is, how does Visual Basic manage complex, multi-tiered distributed applications? Visual Basic uses a Component Manager to keep track of OLE servers on the network. In addition to the Component Manager, Visual Basic uses a Pool Manager that lets you build pools of OLE objects and distribute them to clients. This Pool Manager exists to optimize network performance by helping you avoid the bottlenecks that you may run across when managing so many complex objects over a network. In addition, Visual Basic provides an Automation Manager for managing remote procedure calls to OLE servers.

JET Power

Visual Basic proves that it's ready for your client/server database application by providing several new features. The most significant features include Data Access Objects (DAOs) and the JET Database Engine version 3.0, as well as several handy database tools such as enhanced data control, new bound controls (including a new data-aware grid control), a database forms generator, and an improved database management utility. Visual Basic's Database Manager provides complete control over any database to which Visual Basic can connect. (See Figure 3.)

The heart of Visual Basic's database capabilities is JET. JET's database definition features support the creation, modification, and deletion of tables, indexes, and queries. Also, JET provides attribute-level and record-level data validation. To manipulate data, JET supports SQL, but the SQL is not fully ANSI-compliant.

Visual Basic and JET coordinate database activities using DAOs. These objects enable developers to manipulate information in the database by letting them set properties and access methods encapsulated inside of objects. The DAO 2.0 object model is revamped from Visual Basic 3.0's DAO. New features include multiple workspaces, generic Recordsets, security and referential-integrity objects, property collections, and ODBC SQL pass-through features.

Security and referential integrity in the previous release of Visual Basic were available only if you set them up using Microsoft Access. Visual Basic 4.0 now provides all of these feature through DAOs. The Groups and Users collections of the Workspace objects allow maintenance of the security objects. Also, the database object's Relations collection provides one-to-one and one-to-many referential-integrity constraints, including cascading updates and deletes.

JET 3.0's data replication features are outstanding. JET can make entire copies of a database and reconcile data and schema differences on the fly. Developers need only invoke the MakeReplica method to replicate the database. By invoking the Synchronize method, developers can reconcile an original database with the replicated database. Synchronization can occur between any supported database files, or any database servers using ODBC. Also, JET allows the entire database to be password protected. Prior to JET 3.0, developers could restrict access only to those objects inside the database by granting or denying access to specific users.

But JET 3.0 is not just what you can see, it's also what you can't see. There are several performance enhancements, including a new index structure designed to take up less space and enable faster index creation when keys occur several times in a single table. A new algorithm recycles the deleted pages' IDs, which speeds things up significantly when users are deleting a lot of data. JET can also delete data faster. Also, when compacting the database, JET will actually sort data physically on their primary keys. Thus, the primary key becomes a virtual cluster index.

Other performance enhancements include JET's new-found ability to employ the multithreading capabilities of Windows 95 and Windows NT. JET can also perform database transactions internally to speed-up processing, and employ many

No Other Client/Server Tool Manages Business Change and Complexity Like Magic

Winner — 1995 Droege Developers Competition



Integrating the Future



Managing the Present



Leveraging the Past

The development path you choose determines your company's success. With business practices changing and business environments being reengineered, you need client *and* server application development that stands the test of time. And that can only mean Magic. Magic leverages your legacy IT resources to provide an economical migration path to client/server. It manages your present development needs with an innovative programming paradigm that ensures extremely rapid application delivery and easy maintenance. And Magic integrates the future with fully scalable open architecture that takes applications from tactical to strategic - incorporating new platforms, DBMSs, multiple client/server models and application partitioning scenarios.

The secret of Magic's success is the Magic Engine, an advanced repository of sophisticated application functionality. Driven by a high level, table-driven interface, the Magic Engine lets the developer concentrate on business rules and logic instead of low-level repetitive tasks.

Isn't it time you put Magic to work for your business? For more information and to attend a free Magic seminar, call:

1-800-345-6244

MAGIC

Client and Server Development for Managing Change and Complexity

U.S. Tel: (714) 250-1718, Fax: (714) 250-7404 France Tel: (1) 47 73 90 90, Fax: (1) 47 76 44 36 The Netherlands Tel: 3405 71888, Fax: 3405 71766 U.K. Tel: 01344 303045, Fax: 01344 303024
Germany Tel: (089) 74120-135, Fax: (089) 74120-102 International Tel: (972-3) 538 9292, Fax: (972-3) 538 9333

CIRCLE READER SERVICE NUMBER 40

DBMS/BG

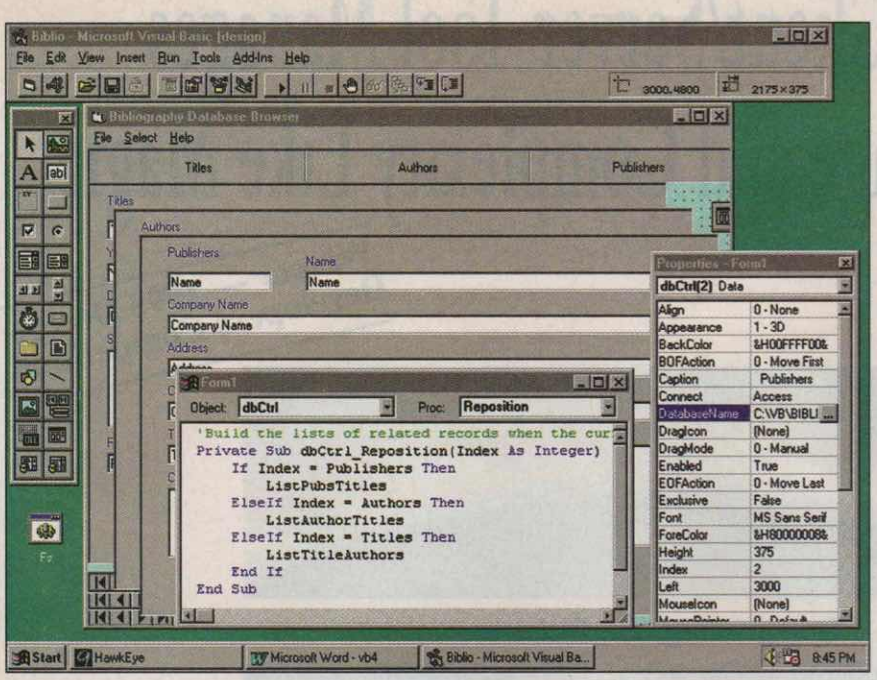


FIGURE 4 **The Visual Basic 4.0 IDE is well designed and lets developers build most of the application visually. The Toolbox is extensible, letting developers add custom controls (OCXs or VBXs) at will.**



FIGURE 5 **Some properties can be set at design time. To set these properties, developers can use the Properties window or property pages.**

sort algorithms now found in the well-tested FoxPro engine. Finally, JET's data cache is limited only by the amount of available resources (memory and disk).

If you spring for the Enterprise Edition, Visual Basic provides a Remote Data Object (RDO) and a Remote Data Control (RDC). These features let you bypass JET 3.0 to take advantage of engine-driven cursors on the server, thus speeding up data access and minimizing network traffic. If you need a reporting tool, Visual Basic bundles Crystal Reports (from Crystal Services, a Seagate Software Company, Vancouver, British Columbia, Canada).

New VBA

VBA (Visual Basic for Applications), the language that drives Visual Basic, is much improved with 4.0. For those of you already programming in VBA, the changes are substantial, and it's not a quick read to get up to speed.

Visual Basic 4.0 is backward-compatible, and you can use the Visual Basic 3.0 style to build applications. However, those of you looking to use Visual Basic 4.0 to its maximum advantage must unlearn the workarounds that were required in 3.0 in order to take full advantage of the new VBA features of 4.0.

For example, Visual Basic 3.0 offered no simple method of passing data from one form to another. Many developers, using workarounds, resorted to using invisible labels to hold shared data, tag properties, or global variables. When using Visual Basic 4.0, forms are objects that developers may extend with custom properties and new methods. Visual Basic developers can add public variables easily, which are referable from any other form in the application. In addition to sharing data from form to form, developers may add custom methods to forms, and these methods can receive parameters passed from other forms.

For those of you starting from scratch, you'll be happy to know that many of the drawbacks in earlier versions of VBA are history with Visual Basic 4.0. The language is fully functional for all application development requirements, and I could not find significant limitations. However, I would not call it a 4GL. With its roots planted firmly in the Basic language, it appears somewhat cryptic. And it will take you some time to get up to speed with VBA. In many respects, it is similar to its main competitor, Borland's Delphi. Delphi is also based on the Object Pascal 3GL.

Building Stuff

For those of you who want to build simple client/server applications quickly, you can't do much better than Visual Basic 4.0. Building a Visual Basic application is as easy as it was with the previous version. The Visual Basic 4.0 IDE is well designed and lets developers build most of the application visually using a development method that I call "drag, drop, and define." (See Figure 4.)

The first step in building a Visual Basic application is to design the forms for the application. If you're building a database application, it's better to define the database and the database connection first. However, it's certainly acceptable to build the forms before worrying about the database.

The Toolbox found on the IDE contains the Visual Basic controls. If the Toolbox is closed, you can open it by choosing Toolbox from the View menu. To place a control on a form, you can either double-click the control, or drag and drop the control onto the form. Visual Basic uses several independent applications running concurrently to make up the development environment. This approach is very different from other tools that provide a single monolithic development en-

WHICH ONE HAS THE BEST REPORTING TOOL?



THEY ALL DO.



If Microsoft®, Symantec®, PeopleSoft®, and 110 other software vendors have one thing in common, it's that they all know a great reporting tool when they see one. That's why they all include Crystal Reports in their software programs.

Crystal Reports' OCX and Report Engine DLL allow for fast, seamless integration into any Windows® database application. And now with Crystal Reports Professional 4.5, you get 16- and 32-bit architecture to generate presentation-quality reports for any Windows environment... in any Windows environment.



The number one reporting tool is even better...now with 32-bit technology.

- The first report writer with both 16- and 32-bit versions (fully compatible) for development for any Windows environment: Windows 3.1, 95, NT or Workgroups.
- Improved Microsoft Access® support lets you enhance reports with OLE objects such as pictures, and report off non-parameterized queries.
- New 32-bit OLE Control (OCX) with >80 properties for full Report Engine control, plus Report Engine access via the VBX, or MFC, or Delphi VCL or 80 direct calls.
- The only reporting tool offering 12 fully integrated graph styles with full customization features, plus custom grouping and extensible formula language.

Upgrade to Crystal Reports Professional 4.5 for just $199 – Risk-free!

Call now to see if you qualify for Crystal Reports Professional 4.5 at the special upgrade price of just $199 (or $395 without a bundled version). In addition, Crystal Reports now comes with another feature: a money back guarantee. Try it for 60 days and if you're not completely satisfied, return it for a full refund. Order today – it's time to get the reporting tool of the experts.

CRYSTAL
A Seagate Software Company

Order now. You may qualify for the upgrade price of $199. 1-800-877-2340 ext. Win 04

Other inquiries call 604-681-3435. E-mail: sales@crystalinc.com. Fax on demand 604-681-3450. Web site: http//www.seagate.com/software/crystal/. Ask for Crystal's FREE newsletter. Crystal and Crystal Reports are trademarks of Crystal Computer Services, Inc. Visual Basic is a registered trademark of Microsoft Corporation. All other names are the property of their respective owners.

Aug. 93 · Oct. 95

vironment with MDI child windows.

The power of Visual Basic is that when the Toolbox is open, you can place new controls in it to extend the capabilities of your development environment. The Professional Edition of Visual Basic provides numerous custom controls. You can also purchase custom controls from third parties. There are hundreds of companies that build custom controls for Visual Basic 4.0 (in both VBX and OCX incarnations). For example, some custom controls provide VCR buttons for applications that use video, controls that provide prebuilt calendars, controls that provide an interface to information servers, and so on. Currently, VBX controls dominate the market, but the majority of third-party control vendors are shifting to OCXs.

The second step in building your Visual Basic application is to set the properties for the objects you have created. Some properties can be set up at design time. To set these properties, you can use the Properties window or property pages. (See Figure 5, page 50.) You may also set properties for multiple controls at one time. To do this, simply click and drag the mouse to select multiple controls on the form. The Properties window displays only the properties common to all the controls, and any change you make to a property applies to all controls.

Depending on your application, you may have to place code behind the controls. For simple applications, you may get away without coding. However, most client/server applications require a significant amount of custom programming for database manipulation and interface processing.

To add a menu to your form, click the Menu Editor button on the toolbar. The Menu Editor is an interactive tool, and you can add new commands to existing menus, replace existing menu commands with your own, create new menus and menu bars, and change and delete existing menus and menu bars.

Having done all that, it's time to test your application and create an executable. The executable, of course, is a misnomer. Visual Basic applications still use a p-code interpreter. This is a major drawback when considering similar competing development environments such as Delphi, which can create standalone native .exes. Performance was okay on a 100 MHz Pentium with 16MB of RAM, but seemed very slow on a 33 MHz 486 with 8MB of RAM (both running Windows 95). So, if you want to deploy Visual Basic on low-powered clients, beware.

Moving Forward

Visual Basic 4.0 is a major upgrade, and it incorporates several new features that make client/server development quick and easy. Moreover, it's now 32-bit and provides much more flexibility for application development. If you're already a Visual Basic developer you should strongly consider an upgrade.

However, there are some things to be concerned about with this release, such as performance and VB's new partitioning features. When comparing the performance of Visual Basic 3.0 with 4.0, the interface performance is slightly worse with the new version, for both the development environment and the runtime. This may be due to the facts that Microsoft expects everyone to use Pentiums and that the tool supports several new power-hungry features. But those of you deploying in the real world will know that everyone has not upgraded to faster processors. However, many of the JET enhancements make up for the slow interface processing.

Visual Basic's partitioning capabilities enable developers to distribute processing on many computers, but, for now, Visual Basic supports only Windows-based servers that can run the OLE Automation Servers. High-end Unix server support would be a nice addition. I would also like to see a transaction-processing object provided with Visual Basic 4.0.

I believe that Visual Basic 4.0 will continue to dominate much of the client/server development market. OLE seems to be the direction of the desktop, and Visual Basic 4.0 makes using OLE a cinch. What it comes down to is, where does a 300-pound gorilla sit? Anywhere it wants to. ■

- Microsoft Corp., One Microsoft Way, Redmond, WA 98052-6399; 206-882-8080 or fax 206-936-7329.

What if—Unproven products and untried methods eat up HALF your Client/Server investment?

Companies are spending two to ten times more money on client/server projects than they had anticipated.

We've found that often half the money is wasted—on unproven products and untried methods.

TO FIND OUT MORE about how our subscription reports and consulting services can help you:

Call 914-939-0045
Fax 914-939-0670
E-mail 72274.1511@compuserve.com

Let our NEW RETAINER CONSULTING SERVICE help you make the MOST of your IT dollars

Then Atre®

ATRE'S Reports from the Trenches—Database Strategies

SUBSCRIPTIONS are available at multiple levels of service, ranging from reports, e-mail and phone consulting to personal on-site consulting by *Shaku Atre.*

Atre Associates, Inc., 222 Grace Church Street, Port Chester, NY 10573

CIRCLE READER SERVICE NUMBER 78

ROGUE WAVE RAVES:

#96: "IF YOU BUILD CLIENT/SERVER APPLICATIONS USING C++, DBTOOLS.H++ IS A MUST."

#97: ". . .DBTOOLS.H++ IS A DREAM COME TRUE. IT PROVIDES EASY ACCESS TO THE MAJOR RELATIONAL DBMSS USING THE SPEED AND EFFICIENCY OF NATIVE API CALLS."

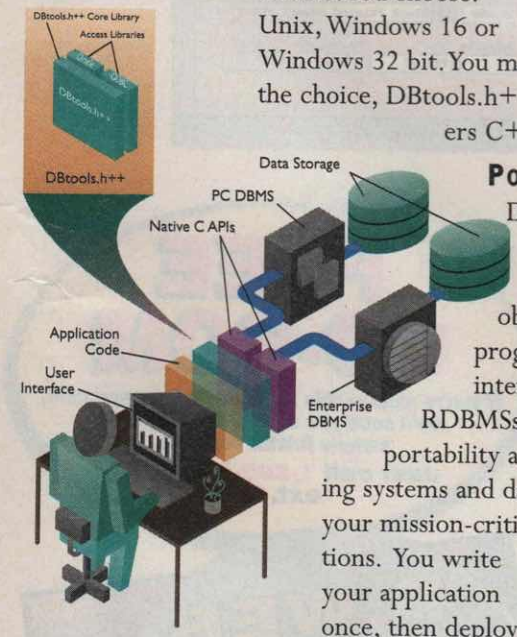
DAVID LINTHICUM, DBMS MAGAZINE, MAY 1995

You Choose the Database. We Deliver the C++ Access.

DBtools.h++

The C++ class library for portable, robust, object-oriented access to relational databases

YOU CHOOSE: Native access to Oracle, Sybase, Informix, or Ingres. You choose: general connectivity through ODBC. You choose: Unix, Windows 16 or Windows 32 bit. You make the choice, DBtools.h++ delivers C++ access.



PORTABLE: DBtools.h++ delivers a common, object-oriented programming interface to major RDBMSs. It provides portability across operating systems and databases for your mission-critical applications. You write your application once, then deploy it to other supported databases without worrying about the differences in data structures or function calls.

OBJECT-ORIENTED: DBtools.h++ delivers the benefits of object-oriented programming to RDBMS developers. You'll write cleaner applications, realize shorter development cycles, and reduce your maintenance costs, without disruptive changes to your IS environment.

FULL-FEATURED: DBtools.h++ gives you the features you need: it encapsulates SQL 92 DML functionality, including critical SQL extensions such as stored procedures. Type safety and data structure compatibility are ensured because your code never tries to access the wrong database. DBtools.h++ also provides a flexible error-handling model that ensures you'll deliver robust production applications.

FREE WHITE PAPER AVAILABLE



CALL US TODAY: Find out how DBtools.h++ can increase your productivity.

Rogue Wave SOFTWARE

Surf Us
http://www.roguewave.com

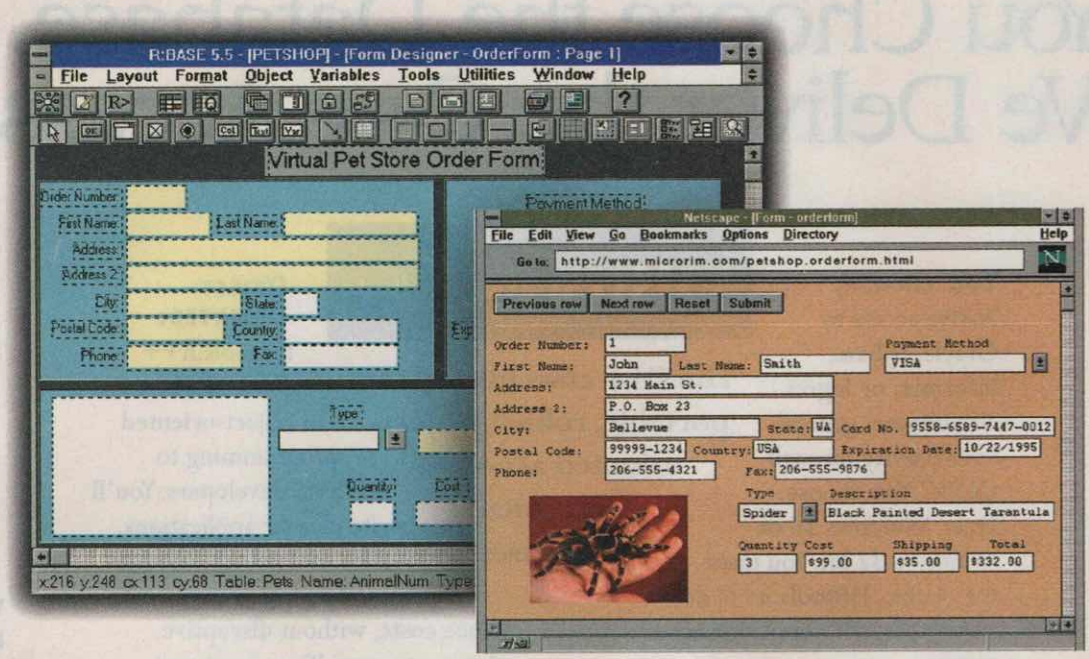
Call Us
(800) 487-3217
or
(541) 754-5010

Fax Us
(541) 757-6650

Email Us
dbtools@roguewave.com

CIRCLE READER SERVICE NUMBER 26

It's so easy we could have called it "Your:WEB"



R:WEB - the first easy-to-use, affordable Internet database solution.



Making a database accessible to visitors of your Web site used to take a good understanding of scripting and a fair chunk of change. Not any longer. Now you can use R:WEB to effortlessly create forms for data exchange between your Web site and your database. R:WEB writes the code - you don't need to know HTML, CGI or PERL scripting. And your database can reside in any ODBC-compliant data source for R:WEB to dynamically link to it. R:WEB maintains security over your sensitive information through data integrity rules, constraints and password security - you have complete control over the information a Web browser can or cannot access. Best of all, R:WEB is built on the powerful R:BASE engine, one of the most respected database engines available. And, oh yeah, did we forget to mention? It's affordable.

So make your Web site work harder the easy way. With R:WEB. And check out the FREE CD-ROM offer.

http://www.microrim.com

FREE CD-ROM

If you're responsible for Web site management, we'll send you a FREE CD-ROM with sample R:WEB software.

Just call 1-800-628-6990, ext.123



MICRORIM®

A Subsidiary of Abacus Software Group
15395 SE 30th Place, Bellevue, WA 98007
(206) 649-9500

R:WEB available through your local reseller or distributor.

Copyright © 1995 Microrim, Inc., a subsidiary of Abacus Software Group. All rights reserved. R:BASE, R:WEB and Microrim are trademarks and registered trademarks of Microrim. Microrim reserves the right to change product and services offered at any time without prior notice.

CIRCLE READER SERVICE NUMBER 202

DBMS-Oriented INTERACTIVE VOICE RESPONSE PART 2

BY P.L. OLYMPIA, PH.D.

PART 2 OF THIS SERIES ON TELEPHONY AND IVR EXPLORES SOME ALTERNATIVES FOR BUILDING COMPUTER-TELEPHONE APPLICATIONS.

Did you know that the first computer-telephone integration (CTI) device is as old as the telephone itself, and not at all a recent invention. It was (and is) called the "teenager." The teenager is really a telephone with hair that can deliver messages at a rate of a million words per minute, and has a computer-like memory for trivia such as who appeared on every MTV program, and when, since the beginning of the network. All I know is that my own interest in CTI came about while watching, with utter amazement, my 13-year-old, slouched on a sofa, carry on two simultaneous conversations over two telephone lines without apparently missing a beat. Between my own CTI work and the teenager, I find that six telephone lines in the house is not enough.

In the first article of this series ("Computer Telephony and Interactive Voice Response," November 1995, page 72), I introduced you to computer telephony (CT) and interactive voice response (IVR). In that article, I gave you some guidelines on IVR script writing and discussed how you could get your first IVR application up and running on a shoestring budget. The Computerfone device I described had three major advantages: low price, transportability, and ease of installation. The last two advantages result from the fact that the device is an external unit that simply plugs into a computer's RS232 port.

Alas, in IVR, as in life, nothing is perfect. If your system requires multiple lines and you don't want to be bothered with programming to the bare metal, you would probably want to look elsewhere. In this article, I discuss some alternative options.

The Periphonics VPS

What would you consider to be the ultimate IVR application? How about an application that allows all of us peasants to file our income tax returns by telephone? My favorite client — the U.S. Internal Revenue Service, for whom my company developed two substantial Visual C++ client/server applications — thinks that's the way to go. No kidding. Can you picture it now? You sign your returns by speaking your own name. Hmm...no more forgery, but that may also mean lifetime employment for folks like Dana Carvey and Rich Little. You'll meet several strange people

.....

Pete Olympia has a doctorate in chemical physics and is the author of several books on dBASE and FoxPro. He is the author of MemoPlus, a memo file handling/repair utility, FPNet, a FoxPro API library of NetWare functions, and RDsecure, a dynamic NetWare rights server that ensures the security of application data.

at a party with worn fingers and no fingerprints. Right away you know for whom they work: H&R Block.

The IRS project is called TeleFile, and it is designed "to make the tax filing process easier and more efficient." Uh-huh. For this large project, the IRS chose the VPS/sp line from Periphonics Corp. VPS/sp is a RISC-based system hosted on a Sun SPARC machine that can support more than 120 digital or analog telephone lines. To complement this hardware, Periphonics offers a visual application development tool for designing and creating IVR applications. The tool, called PeriProducer, lets users build applications by connecting iconized building blocks. The building blocks include not just the bread-and-butter functions of an IVR application such as "Answer the phone with a greeting" and "Get caller data" (touch-tone or voice), but also extended functions such as "Send fax" or "Receive fax." Thus, with this tool, developing an IVR application with integrated fax-on-demand is almost a no-brainer.

Just to give you a flavor for some of the IVR applications developed by government and large businesses using the VPS/sp line, the California Department of Motor Vehicles has a system in place to handle motorists' inquiries around the clock, and also to enable motorists to schedule appointments. The system probably got its most severe shakedown the day a white Bronco was seen leading a parade of police cars down the Hollywood Freeway.

One of the best examples of IVR's ability to improve service and cut costs is the system operated by the Army Emergency Relief (AER). Soldiers who are transferring from one installation to another must first clear any debt with the AER headquarters in Virginia. As I understand it, there are approximately 20,000 soldiers in transit around the world at any given time. In the past, considering the volume and the 11 time zones involved, an installation that could not reach headquarters tended simply to clear and release soldiers even though these soldiers may have had outstanding debts. But this isn't the case anymore. The VPS/sp IVR application speeded up the clearing process with its 24-hour availability, and it reduced the problem of outstanding AER accounts substantially.

But what if your organization or client does not have the budget of the AER or IRS? You may not be able to play in the VPS/sp league, but you're not out of the IVR market altogether if you talk either Microsoft FoxPro or Visual Basic. Read on.

The VoysAccess System

The IVR market has been exploding during the last few years, and this explosion has resulted in a wide variety of application development choices. The burning question is: How do you choose the right package? I don't want to get into a religious war here, but my choice of a development package always has been based on a DBMS with a robust programming language, rather than on a conventional programming language or a specialized IVR code generator. From where I sit, and given the small budget of typical clients, I consider FoxPro to be a good choice.

One of the best examples of IVR's ability to improve service and cut costs is the system operated by the Army Emergency Relief.

FoxPro aficionados have waited a long time for the opportunity to play in the IVR game. Therefore, they were dancing in the streets the day Voysys Corp. released the VoysAccess toolkit (VTK). The toolkit allows a reasonably accomplished FoxPro for Windows developer to design and implement an IVR application in a matter of hours. VTK consists of a two-line telephone interface board made by industry-leading Dialogic Corp. (Parsippany, NJ), IVR development software in the form of a FoxPro/Windows API library, a set of pre-recorded voice prompts, and VoysSmith, a voice prompt editor.

Once you get past the installation, you're home free. The installation's level of difficulty increases in inverse proportion to your experience in mounting interface boards, as well as to old-fashioned luck. The board requires: 1) a hardware interrupt (default IRQ3); 2) a shared memory address (default D000H) used for communication between the board and the software; and 3) a software interrupt (default 62H). The first two pieces require that you change jumpers on the board if you can't live with the default values.

I used IRQ 5 (typically used by LPT2) because my machines already have IRQ 3 and 4, taken by a mouse and modem, respectively. The installation software does not help you figure out which IRQs are taken; instead, you are asked to run the MSD (Microsoft Diagnostics program) utility to get that information. Windows 95 (a.k.a. Macintosh 1988), where were you when I needed you? If you are using an expanded memory manager such as EMM386 or Quarterdeck's QEMM, you'll need to exclude 200H bytes of memory beginning with your chosen shared-memory address.

The install program adds a line to your autoexec.bat file, invoking the batch file Voysload.bat. You must run the latter before you can use the VTK. It loads a memory resident module that grabs a whopping 150KB of memory. You must also run the VoysAccess Server (click on the icon or make it part of your Windows startup) before you can run the VTK. The memory resident module and the Server are required before you can "talk" to the voice board, even at runtime.

Using the Toolkit

VTK comes with a few ready-to-run examples. A good one to experiment with is ParPhone.prg, which is an application that allows part-time workers to call in, check their work assignments, and accept or decline one or more assignments. I must warn you that the sample programs are not exactly world-class FoxPro program code: They work, but they simply show how to use the VTK functions, not to teach users FoxPro programming.

You may be surprised to learn that the VTK has only eight functions. The truth is, you don't even need all of those functions to develop fairly elaborate IVR applications. Thus, it's a small wonder that you can get going in a matter of hours. Basically, there are two kinds of IVR applications: one that people dial in to call, and one that dials out to call people. Most applications I have seen are of the first type. In such applications, the program flow is predictable, consisting of the following basic steps:

1. Wait for the phone to ring, then pick it up.
2. Speak a greeting or prompt for caller input.
3. Get caller touch-tone input.
4. Perform database action based on caller response.
5. Repeat steps 2–4 until input ceases.
6. Hang up the phone and go back to step 1.

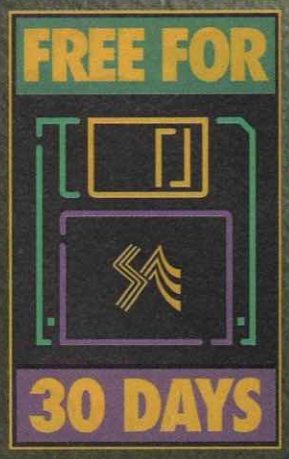
You don't have to be a Nobel laureate to figure out that the heart of the application requires only four IVR functions: VWaitRing(), VSpeak(), VGetTones(), and VHangUp(). I bet you also figured out that these functions are used in steps 1, 2, 3, and 6, respectively. See? You already know how to write an IVR application, and you're not even a brain surgeon. Of course, you want to check periodically to see if the caller has hung up on your application, so you'll need a VChkHangup() function. If you allow users to leave a voice message, you'll also need a VRecord() function.

Time and Attendance IVR Application

I developed a complete, commercial-quality, employee time-reporting IVR system in less than a week — not because I walk on water, but because I used two superb FoxPro-based tools: the VoysAccess tool-

(continued on page 76)

We're Free for the next 30 days, if you are.



System Architect is the award winning client/server software engineering tool that delivers higher performance — at a lower price. And to show you, we'll send you a copy of System Architect™ to evaluate for 30 days — absolutely free.

Object-Oriented A&D

Implement your OO design using SA Object. Some of the features include:

- Support for Booch, Shlaer/Mellor, OMT, Coad/Yourdon, and Use Case
- User-extensible repository
- Automatic creation of data models from class models
- Roundtrip engineering with C++
- Generates Smalltalk

Advanced Data Modeling

Develop your server design using SA's data modeling capabilities:

- Comprehensive support for the DBA function
- Generates DDL from data models for Oracle, Informix, SQL Server, Access, SYBASE and many others
- Support for physical/logical data modeling
- Specify referential integrity rules
- Trigger & Stored Procedure editors



SA's family of tools stores BPR, Structured Analysis & Design, Object-Oriented Analysis & Design and Data Modeling information in a multi-user, shared repository.

Structured A&D

Develop your client/server design using:

- Information Engineering
- Yourdon/DeMarco (event-driven)
- Gane & Sarson
- Ward & Mellor (real-time)
- Decomposition diagrams
- SSADM IV

BPR

Popkin Software has introduced a business process reengineering tool. Call for details.

55,000 Users Can't Be Wrong

That's the number of satisfied SA users worldwide, and it's growing every day.

Try It Free For 30 Days.*

To Qualify For An Evaluation Copy Of SA, Call 800-732-5227, ext. 698 or Fax 212-571-3436.

* Offer good in US and Canada only. See distributor telephone numbers below for international response.

SYSTEM ARCHITECT™

Popkin Software & Systems, Inc.
11 Park Place, New York, NY 10007
Real Tools For The Real World.

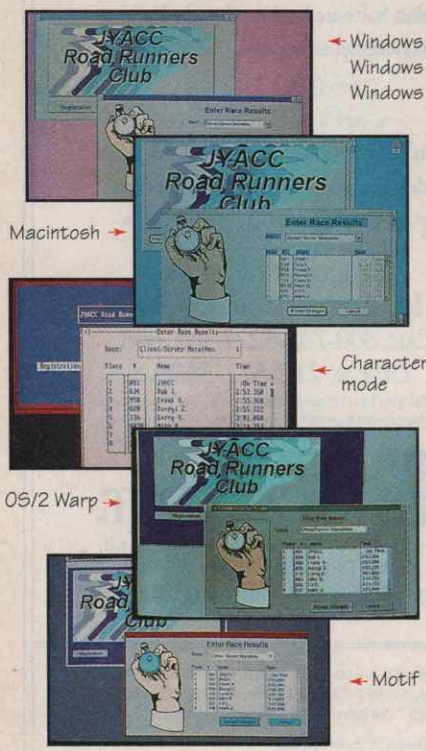
Australia (1-800) 658660 * Benelux +31-30-666 55 30 * Brazil 55-11-546-0246 * Chile 56-2-334-1869 * Colombia 57-1-410-1303 * Denmark 45-45-999300
England 44-1926-450858 * Germany 49-6151-866-620 * Iceland 354-5-671920 * Italy 39-49-8700366 * Japan 81-462-73-5922 * Korea 212-751-1109
Malaysia 60-3-757-1806 * New Zealand 61-02-346499 * Spain 34-3-415-7800 * Sweden 46-8-626-8100 * Switzerland 41-61-6922-666

©1996 Popkin Software & Systems, Inc. The System Architect logo is a trademark of Popkin Software & Systems, Inc.
All other brand and product names are trademarks or registered trademarks of their respective holders.

CIRCLE READER SERVICE NUMBER 118



Building Client/Server Applications Doesn't Have to be Painful – Get JAM® 7 a Smashing New Version!



Whether you are a business analyst or a top seeded developer - JAM 7 makes it easy to develop and deploy industrial-strength applications. The newest release of JYACC's industry leading second generation client/server tool gives you the features and flexibility you need to see complex projects successfully through from start to finish. JAM is a scalable solution that spans single-tier through three-tier architectures; desktop through mainframe platforms; departmental through enterprise-wide systems; and individual through large development teams. So don't bear the pain of inadequate tools - get JAM and be a winner!

Empower your team with the JAM advantage:

- ✓ Portability to Windows, Motif, Mac, OS/2 Warp and Character-mode
- ✓ Repository-Driven Development
- ✓ Full Support for Native Database Features
- ✓ Graphical Object-Based Environment
- ✓ Database Screen Wizard with Full Transaction Control
- ✓ Automatic SQL Generation
- ✓ Advanced Business Graphics
- ✓ Team Development Support
- ✓ No Runtime Charges

You Make The Call: 1-800-458-3313

or E-mail: jam7@jyacc.com for a free **demonstration kit.**
Visit our Web site at http://www.jyacc.com
For international inquiries call: 1-212-267-7722 or FAX 1-212-608-6753

BRAZIL (55) 11 816 6229 • DENMARK (45) 33 32 55 77 • FINLAND (358) 0 162 966 • FRANCE (33) 1 46 92 45 44 • GERMANY (49) 40 79 70 07 0 • HRVATSKA (385) 1 242 116
INDIA (91) 22 283 1188 • ISRAEL (972) 3 921 8320 • ITALY (39) 2 45 2701 • MEXICO (52) 5 550 4500 • RUSSIA (7) 095 288 1924 • SAUDI ARABIA (9662) 671 8749
SINGAPORE (65) 220 8322 • SLOVENIA (386) 61 1405 004 • SPAIN (34) 1 804 0625 • SWEDEN (46) 8 96 10 10 • SWITZERLAND (41) 21 991 9041 • THAILAND (66) 2 513 3559
THE NETHERLANDS (31) 70 320 9214 • UNITED KINGDOM (44) 171 814 6660

JAM is a registered trademark of JYACC, Inc. Other trademarks are the property of their respective owners.

JYACC
Meeting the needs of professional developers for over 16 years

CIRCLE READER SERVICE NUMBER 139

BUILDING A DBA REPOSITORY SYSTEM

APPLYING APPLICATION-CENTRIC REPOSITORY PRINCIPLES TO PHYSICAL DATABASE MANAGEMENT.

BY MARTIN RENNHACKKAMP

In this two-part series, I explain how you can use repository technology for physical database management. In this first part, I provide an overview of repository technology and the challenges of physical database management. I also discuss the shortcomings of existing repository-based tools with respect to physical database management, and describe how repository technology is applicable to physical database management. Next month, I will describe a repository-based system that my company has developed for physical database management.

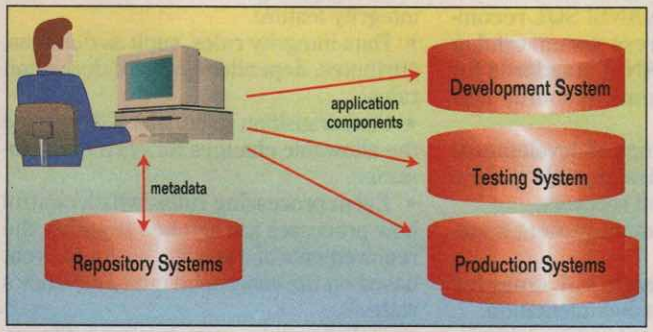


ILLUSTRATION BY BARBARA LICATA

FIGURE 1 **A repository system. The repository is a knowledge base that contains metadata about the information an enterprise keeps and accesses, the processes it performs to be successful, and how it has implemented its application systems to support its information processing needs.**

Repository Technology

A repository system's purpose is to record and manage the information about the data models and the designs and implementations of the enterprise databases, applications, and systems. In other words, the repository system records and manages the metadata of the enterprise. The repository is a knowledge base that contains metadata about the information an enterprise keeps and accesses, the processes it performs to be successful, and how it has implemented its application systems to support its information processing needs. A repository system is illustrated in Figure 1.

Thus, a repository system is a database-centric application that manages corporate metadata and therefore its data and application systems. The result is an integrated and manageable corporate information resource. Repository technology is a key component in the integration of the enterprise's information environments. It is the technology used to inte-

.....

Martin Rennhackkamp is the owner and principal consultant of The Data Base Approach, a corporation specializing in relational and distributed databases, based in Cape Town, South Africa. You can reach Martin via the Internet at mr@dba.co.za.

grate different data models, independent databases, diverse and heterogeneous database technologies, user applications, CASE tools, reference and directory management systems, systems administration, and object management services. The repository system presents a single, consistent view of the metadata describing all of these pieces.

Content

The metadata stored in the repository system typically includes the following:

- Data dictionary metadata includes relations (such as tables), keys, attribute details, formats, integrity constraints, and business rules applied to database structures.
- Design metadata includes requirements documents, design documents, architectural diagrams, data flow diagrams, structure charts, hierarchy charts, and project plans. This includes logical and physical data models, business rules (for example, obtained from forward- or reverse-engineering), and integrity rules.
- Application development metadata includes version-control data, source code, test scripts, and quality assurance reports (including physical data structures for DBMSs and 4GLs), as well as query documentation for end users.

The repository system must not only track metadata such as physical data elements, screens, forms, and tables, but it must also track business rules, in which some data elements might be related to other data elements. For example, it must track which data elements appear on which forms, screens, and tables. Therefore, the information about the relationships between the data elements and the screens should also be maintained. The metamodel must contain all of the valid associations, in addition to the objects and properties that make up the metamodel.

Facilities

The repository system's facilities are implemented as database applications operating on normalized sets of databases, object bases, knowledge bases, and files to manage the metadata at all levels of abstraction. These facilities, which provide standardized services for data access, data sharing, and data management, include:

- Modeling tools (that is, CASE tools) for recording and analyzing enterprise knowledge in the form of data and process models.
- Library facilities, such as dictionaries, encyclopedias, thesauruses, and glossaries.
- Directory management facilities to record data addresses and attributes, in order to provide naming services.
- Systems administration facilities to install and manage information databases.
- Repository management facilities, such as object management services, time services, and interfaces to other services.
- Interactive interfaces with DBMSs, 4GLs, 3GLs, precompilers, and other software components, such as scheduling packages, all of which provide the ability to cross-reference and synchronize jobs, programs, and data definitions.

Functions

A repository system must support the storage, retrieval, version management, and configuration control of all the metadata used for the many aspects of application development. It must provide the user applications with a consistent scheme for managing data in a distributed information processing environment. The application processes must be able to share data across heterogeneous computer networks, yet remain insulated from the details of interaction with multiple types of hardware and software. The repository system must provide a uniform services interface to application processes, and it must maintain knowledge of the physical location of the data, as well as control the distribution thereof.

A repository system must provide the following functions:

- Support and encourage repository sharing resources across multiple business units throughout the enterprise.
- Provide concurrent, multiuser access.
- Provide a robust reporting environment to users and technical specialists.
- Store multiple process and data models of business units with multiple versions of each.
- Support a metamodel that is adaptable to change and enforces data integrity constraints.
- Provide a solution consistent with existing frameworks and standards, such as IBM's Repository Manager framework, John Zachman's Information Systems Architecture (ISA) framework, and various ANSI standards (for example, the ANSI/X3/SPARC three-schema architecture, some of the ANSI SQL recommendations that address system catalog interfaces, and the ANSI standards for an Information Systems Resource Dictionary [IRDS]).
- Maintain descriptions of development objects for analysis, design, code generation, help screens, and recovery.
- Maintain change-management histories.
- Maintain and generate 4GL cross-references and program documentation.
- Maintain component engineering information.
- Provide the ability to add object types and define their interaction (objects such as programs, systems, jobs, files, fields, and reports integrate with each other).
- Provide the ability to add indexing words to objects for searching.
- Know which repository facilities are running and accessible on the network.
- Know what data resides on the network, how to retrieve it, and whether the specified data is currently retrievable.
- Enforce security constraints on network-based repository resources on the network.
- Provide Code generation.
- Provide Impact analysis.

Some code, such as data definition statements, can be automatically generated from the information stored in the repository — as long as the repository and the DBMS are tightly integrated. The input for physical design is in the repository (definitions, names, lengths, and formats). With these elements, you can create constructs for testing and later move them to production. Because database definitions are usually DBMS-specific, the DBA will have much less work to do if the repository can address the DBMS in its own terms (for example if the repository can generate the data definitions in a syntax supported by the DBMS, or, even better, if it can run the database definitions directly against the databases managed by the DBMS).

The repository's role is to identify the impact made when programs alter the way they use data or when the data structures change. An analyst must be able to determine the impact of these changes, even if the only resource is a list of potentially affected relationships. Cross-referencing is good for more than just impact analysis. Using even a partial cross-reference with the actual program descriptions results in "asset management"; that is, the ability to increase the reusability of subroutines or other program code. Impact analysis should also examine the nature of the following business rules affected by the change:

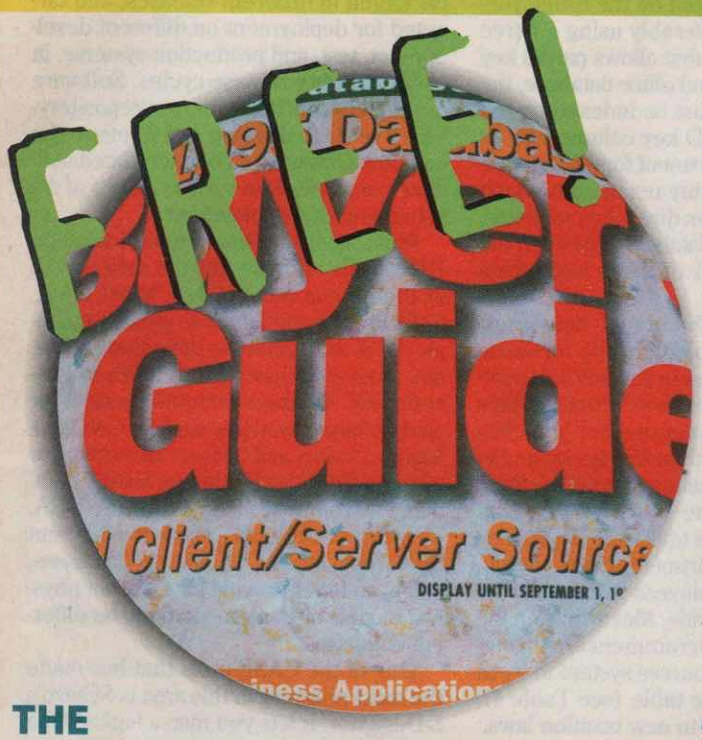
- Entity instance relationships, typically enforced through the DBMS's referential integrity feature.
- Data integrity rules, such as domains, attributes, dependencies, and derivation rules.
- State transition rules, which determine the allowable changes for each entity instance.
- Event processing rules, which control how processes are invoked to satisfy the requirements of a specific business event based on the associated entity instance's state.
- Delivery rules, which govern the external interfaces such as screens, dialogs, and reports.

In multi-release impact analysis, you must perform the same impact analysis

Check it Out!

Oracle Workgroup 2000, Powersoft's S-Designor, the latest scoop on IBM's DB2 Version 2, demo tools by Popkin, Btrieve, JYACC, Cast & Intersolv

ALL ON YOUR LUNCH HOUR!



Major Sponsorship for this CD-ROM Provided by:

Oracle Workgroup 2000™ Proof that you can.

IBM®

THE 1996 *DBMS* *BUYER'S GUIDE* ON CD-ROM— BETTER THAN EVER

ALL NEW VENDOR DEMOS!

- **Oracle:** Oracle7 Workgroup Server for networks, Personal Oracle 7 desktop DBMS, Oracle Power Objects — all three fully-functional 90-day evaluation versions
- **IBM:** multimedia presentation about DB2 Version 2
- **Powersoft:** evaluation versions of the S-Designor Professional 4.2 CASE tool, and Watcom SQL Server DBMS
- **Intersolv:** multimedia presentation and white paper on the DataDirect suite of ODBC connectivity, query and analysis tools
- **JYACC:** Windows demo of JAM 7 cross-platform application development tool
- **Popkin:** multimedia presentation and spec sheets on System Architect 3.1 CASE tool
- **Cast Software:** presentation on CAST Workbench administration tool for SQL Server
- **Btrieve:** presentation on Scalable SQL DBMS

ALL NEW USER SOFTWARE!

- NETCOM NetCruiser Internet service and Web browser. Use this to view our CD-ROM product listings right off the World Wide Web! You can access the Internet at a special low price of just $9.95 for the first month.
- NavCIS CompuServe navigator from Dvorak Development, preconfigured for the *DBMS* Forum on CompuServe
- U.S. five-digit zip code database courtesy of Alpha Software Corp. to use for your database applications
- Star Tracker, a data warehouse reporting tool written by *DBMS* columnist Ralph Kimball
- Viewing software for documents in these applications: Microsoft Word and PowerPoint, Microsoft Video clips, and Adobe Acrobat.

AND DON'T FORGET THE PRODUCTS!

Now you can instantly find the database and application development tools you need from among our 800+ detailed product listings from 400 vendors — everything from tools to servers to middleware. With our user-friendly interface, you'll find that special product in seconds!

YOU COULD WIN UP TO $500!

A limited number of our CD-ROM Buyer's Guides will have a secret key in the advertising section that, when activated, could spell M-O-N-E-Y for you! SO WHAT ARE YOU WAITING FOR? Grab that CD right now, and put it in that drive. Our CD-ROM Buyer's Guide will save you valuable time and money.

LET US KNOW WHAT YOU THINK

If you've got any comments, suggestions or feedback for us, just send it to our Technical Editor, Maurice Frank via the Internet at: mfrank@mfi.com.

ORDER EXTRA COPIES NOW

Send a check or money order payable to Miller Freeman, Inc., to:

***DBMS* Magazine**
411 Borel Avenue, Suite 100
San Mateo, CA 94402
ATTN: Gae Reck

DBMS

Domestic orders are $19.95 each, $24.95 each for all orders outside the continental U.S. Don't forget to include your name and address.

Check out these vendor demos!

BTRIEVE TECHNOLOGIES · CAST SOFTWARE · INTERSOLV · JYACC · Powersoft *Building on the power of people.* · SYSTEM ARCHITECT POPKIN SOFTWARE & SYSTEMS, INC.

TABLE 1. The Employee table

EMPLOYEE		
EmployeeID	EmployeeName	Department
1345	John Walker	Sales
4692	Sue Brown	Repairs
1467	Pam Johnson	R&D

TABLE 2. The EmployeeLeave table

EMPLOYEELEAVE		
EmployeeID	FromDate	ToDate
1345	1 Jan 1995	15 Jan 1995
1345	5 Apr 1995	10 Apr 1995
4692	1 Jan 1995	8 Jan 1995

for each application version in order to determine which application components have been affected by a change. Another change-management challenge occurs when the release schedule is restructured. When this occurs, you must conduct careful impact analysis.

Physical Database Management

Database management, in the generic sense, covers the entire life cycle of all of an organization's databases, from requirements definition through logical database design to physical database design, implementation, and performance tuning on various target databases. Logical database design is performed by a data administrator. This administrator documents the logical data model, typically using a CASE tool. The goal is to represent the users' business needs as a consistent and correct logical data model. The data administrator conducts JAD sessions with users to establish a global logical data model, which is sometimes partitioned for particular application areas.

A DBA performs physical database management. Physical database management carries the logical data model forward to physical design. The goal of physical database management is to implement and manage the logical data model as an efficient, consistent, and correct physical data model.

A large organization typically has a number of databases, often in various stages of development (for example, regional production databases in its regional offices, a central production database in its head office, an evolving warehouse database for statistical processing, and various test and development databases). Developers must manage the implementation of the physical data model on the various databases of the enterprise.

Obviously, physical database management is complex. DBAs must manage multiple physical implementations of the same or different logical data models in different databases, and concentrate on tuning various aspects of these different databases to improve performance. In addition to the widely diverging performance tuning aspects, DBAs must also manage various other implementation-related objects, such as triggers, user-defined procedures, rules, events, and views.

As an example, see Tables 1 and 2. In the sales branch's database, the Employee table must be indexed on the EmployeeName column, preferably using a b-tree or ISAM structure that allows partial key searches. In the head-office database, the Employee table must be indexed on the unique EmployeeID key column and on the repeating Department foreign-key column, both preferably using some hash structure that allows direct key look-ups. In the warehouse database, however, the view (see Table 3) of these two tables would be used.

In addition, on the branch databases, the DBA must implement the business rule that employees in the sales department may not take more than 30 days leave, while employees in other branches may not take more than 25 days leave. On the head-office database, the DBA must also implement the business rule that there may only be a total of 25 employees in all the sales departments country-wide, and only eight employees in the R&D department country-wide. Meanwhile, in the IT section, the programmers are changing the human resources system to work with the Employee table (see Table 4), which would cater to new taxation laws.

Available Tools

The database management function must manage complex, interrelated objects such as table definitions, integrity constraint definitions, business rules, views, stored procedures, and performance improvement facilities. In an environment with multiple DBMSs, this is even more complex because each vendor's implementation of these database objects differs. For example, CA-OpenIngres implements business rules through stored rules and procedures, while Oracle and Sybase implement business rules through triggers. However, Oracle's and Sybase's SQL syntax triggers differ significantly.

Software developers work in a slightly different environment than that of a DBA. They work with different versions of program code, written in different programming languages (such as C, C++, Cobol, Visual Basic, and various 4GLs), addressing different target systems (such as Windows for DOS and Motif for Unix), packaged in different releases, and targeted for deployment on different development, test, and production systems, in ever-changing release cycles. Software developers can, however, use repository-based tools to develop and manage the complex environment of interrelated software components in various stages of development and deployment.

Most repository-based tools, including CASE tools, address only the subject area of the data administrator. Most CASE tools have the facilities to generate simple ANSI SQL database definition scripts, and most even have the facilities to generate SQL database definition scripts for specific target systems, such as CA-OpenIngres, Oracle, and Sybase. However, few of these CASE tools address the full complex environment managed by the DBA. Very few CASE tools provide for different logical data models on different databases, and even fewer provide for different physical models or implementations on different databases.

One of the CASE tools that has made the most progress in this area is Sybase's S-Designor. It lets you map a logical data model (defined in terms of entities, attributes, and relationships) to various physical data models (defined in terms of tables, columns, and constraints), as supported by different DBMSs (for example SQL table definitions in the SQL dialects supported by Sybase, Oracle, and CA-OpenIngres). For a discussion and a listing of CASE tools, see "The Evolution of Client/Server CASE," on page 67.

TABLE 3. The Employee view of Tables 1 and 2

EMPLOYEE			
EmployeeID	EmployeeName	Year	DaysLeave
1345	John Walker	1995	22
4692	Sue Brown	1995	6
1467	Pam Johnson	1995	0

TABLE 4. The Employee table with taxation laws included

EMPLOYEE				
EmployeeID	EmployeeName	SocialSecurityNumber	Department	AppointmentDate
1345	John Walker	34-89675-234	Sales	4 Jun 1980
4692	Sue Brown	67-11201-829	Repairs	1 Jan 1994
1467	Pam Johnson	45-24011-114	R&D	1 Apr 1992

Most CASE tools will allow multiple versions of the Employee table. This approach is very useful for the evolution of the Employee table used in the branch and head-office databases. Few CASE tools, however, let you associate versions 1 and 3 with the branch and head-office databases, while version 2 is deployed on the warehouse database.

Very few CASE tools provide for physical tuning aspects, such as table structures, fill factors, and re-indexing scripts. For example, few CASE tools make the distinction that: tables can be structured as b-tree, hash, heap, or ISAM structures in CA-OpenIngres, and as b-tree or hash structures in Oracle; you can specify a table preallocation size in CA-OpenIngres; and you must specify a table preallocation size in Oracle. Very few CASE tools address the extended features provided by the more advanced RDBMSs such as CA-OpenIngres, Oracle, and Sybase. These include objects such as procedures, rules, triggers, events, and views. For example, few CASE tools let you construct a view as a SELECT query about existing tables and other views, with full cross-referencing information on which columns are used in which views. Finally, few CASE tools let you capture stored procedure and trigger definitions, let alone address the subtle differences between implementation and syntax (for example, between CA-OpenIngres, Oracle, and Sybase).

A few repository-based CASE tools do address the interface between logical database design and physical database implementation through simple SQL script-generation facilities. However, the database management function must still use the tools and facilities of the operating systems and the DBMSs to manage physical database management in the complex database environment.

Applying Repository Technology

Despite the challenges, it makes perfect sense to utilize repository technology for physical database management. If you consider the *contents* of the repository system, the following applies to physical database management:

- The database objects described by the data dictionary metadata — in other words, the relations (that is, tables), keys, attribute details, formats, integrity constraints, and business rules — are the same objects managed by the DBA. The focus just changes from the objects accessed by the application programs to the primary objects being managed. This includes adding all the implementation and deployment details of all these various objects.
- The design metadata still includes requirements documents, design documents (such as entity-relationship diagrams, relation diagrams, entity-life histories, data fragmentation, and allocation schemas), data flow diagrams (these still represent the program entities accessing the database objects), structure charts, hierarchy charts, and project plans. For the DBA, the focus shifts once again to the physical implementation details (that is, the differences between the operational systems running on CA-OpenIngres and the warehouse system running on Oracle), the physical deployment details (which version of which table is deployed on which database, and which database is implemented using which DBMS), and physical design details (that is, the different indexes used in the branch and head-office databases).
- The application development metadata still includes versions of applications and queries accessing the various database objects, including references to control data, source code, test scripts, and quality assurance reports. The metadata also includes physical database structures implemented in the different DBMSs to cater to the application logic implemented in the various queries, 3GL, and 4GL programs.

When you consider the *functions* provided by a repository system, the following applies to physical database management:

- The repository resources, including all of the database object definitions, are shared across multiple business units throughout the corporation. All of the DBAs in the various projects and departments must share one global logical data model, with its various physical implementation details.
- DBAs, data administrators, and application developers must have concurrent, multiuser access to all the metadata used for physical database management.
- The DBAs, data administrators, and application developers must be able to report on the physical database management metadata from any perspective.
- The repository system must store multiple physical data models and the process specifications used by all business units, maintaining multiple versions of each as they evolve and as they are applied to different physical databases.
- The meta-model used for physical database management, similar to its application-centric counterpart, must be adaptable to change. It must also enforce data integrity constraints to ensure that the metadata used for physical database management is correct and consistent.
- The metadata used for physical database management must also support existing frameworks and standards, such as John Zachman's ISA framework and the ANSI SQL standard.
- Descriptions of all the database objects must be maintained for analysis, design, code generation, help screens, and recovery.
- Change management histories of all the database objects must be maintained to illustrate how they have evolved over time (to provide for roll back to previous releases, for example).
- You must be able to generate 4GL and program cross-references and documentation (such as which tables and views are used in which queries, which stored procedures are called by which applications, and so on).
- The component engineering information used for physical database management is extensive. Tables and views are used in view definitions and procedure definitions, all of the database objects are grouped into schemas, schemas are deployed on logical databases (such as a branch database), and logical databases are implemented as physical databases (such as the branch database in the Miami sales office).
- The ability to add database object types and define their interaction is also important. DBMSs are constantly evolving and constantly supporting more and more diverse types of objects stored in the database. Examples include declarative referential constraints, stored procedures, triggers, and database events. Current extensions provide for objects, persistent objects, query execution plans, optimizer statistics, and distributed database object definitions.
- Similar to its application-centric counterpart, it should be possible to add indexing words to the managed objects and search on the basis of those indexes.
- The repository system must know which database management facilities are running on the network and whether they are currently accessible. The system must also know what data resides on the network, how to retrieve it, and whether the specified data is currently retrievable.
- The repository system must enforce security constraints on the network-based repository resources.
- The repository system must be able to generate code for all the facets of physical database management (such as create, modify, optimize, restructure) of all the database objects (such as domains, tables, views, triggers, events, schemas, and even entire databases) for multiple target systems (such as Oracle, CA-OpenIngres, Sybase, and so on), each with its own subtle syntax differences.
- Impact analysis is very important for physical database management. The DBA must be able to determine the impact of any change to the physical database. For example, if a table's column changes, which indexes, views, procedures, triggers, and events are affected by the change? In which databases is the specific version of the table used?

TABLE 5

ASPECT	APPLICATION-CENTRIC REPOSITORY SYSTEMS	PHYSICAL DATABASE MANAGEMENT REPOSITORY SYSTEMS
Share repository metadata across business units	Application developers share global definitions and application libraries	DBAs share physical database objects and definitions
Multiple versions of models	Versions of application design and implementation documents	Versions of logical and physical database design documents
Multiple versions of object definitions	Versions of program code, shared libraries, test scripts, and so on	Versions of table and view definitions, trigger code, database definitions, and so on
Adaptable/extendable meta-model	Application development approaches change	Object types stored in databases change
Existing frameworks and standards	John Zachman's Information Systems Architecture (ISA) framework; AD/Cycle framework	John Zachman's Information Systems Architecture (ISA) framework; ANSI SQL standard
Descriptions of objects for analysis, design, and implementation	Application models, application designs, program code libraries	Data models, database designs, database definition scripts
Change-management histories	Versions of application designs; versions of program code	Versions of database designs; versions of database definition scripts
Cross-referencing documentation	Interrelated application objects (such as procedures calling other procedures); database objects used by applications	Interrelated database objects (such as tables used in views and stored procedures); application components accessing database objects
Component engineering information	Code in functions in procedures in application systems	Columns in tables in views in stored procedures in schemas in logical databases in physical databases
Code generation	Code libraries, programs, application systems	Database management scripts, database definition scripts
Change/impact analysis	Application components affected by changes in other application components	Application components and database objects affected by changes to other database objects

Comparison of conventional application-centric repository systems and physical database management repository items.

The most important analogies between the functions provided by conventional repository systems and a repository system adapted or extended for physical database management is shown in Table 5.

The proliferation of repositories is a big problem. The "theoretical" correct approach is to use *one* standard, integrated repository for all these functions — which results in totally integrated metadata management. In practice, however, this does not happen (yet), because all these products have their own, fairly proprietary, repositories.

In the DBA Repository System (which I explain in detail next month), we took the viewpoint that most CASE tools can generate SQL-92 syntax (at least some adherence to standards!), with some variations for each particular DBMS. We therefore provide an import facility with which you can "synchronize" the DBA Repository from your CASE tool repository. It doesn't blindly bulk-load the metadata into the repository database; rather, it matches and resolves previous versions of the defined objects where it can. This is how we overcome the workload issue — no one, not even different developers on the same project, wants to capture the metadata twice, because that would leave open the possibility for discrepancies between the logical and physical data models.

Until Next Month...

DBAs and application developers work in equally complex environments. They perform similarly complex tasks in order to manage multiple interrelated database objects for different target systems. DBAs inspect and manage the same managed objects, but from a different perspective — they actually manage the database objects traditionally cross-referenced from applications, and they cross-reference the application objects themselves. However, they do not have highly integrated repository systems at their disposal.

Although some of the repository-based development systems cater to different logical and physical models, and although some of them can help the DBA deploy the physical model on development, testing, and production databases, they do not provide sufficient functionality for the DBA's day-to-day tasks. For example, they do not provide functions for re-indexing tables, re-gathering optimizer statistics, and re-allocating tables to different physical devices. On the other hand, the tools traditionally used by the DBA to perform these tasks, be they from a third party or part of the DBMS itself, are usually not repository based. They only contain the present state of the database objects, without facilities to document or manage previous or future versions of the database objects.

I have tried to show in this article that you can apply the functions performed by a conventional application-centric repository system equally well to physical database management. It's just the focus that changes. The repository objects that DBAs would use most are the actual database objects (that is, the tables and views to be accessed), with the relationships, integrity, and other constraints defined among them, as well as the where-used information (that is, which tables and views are used in which queries and applications).

Next month, I will describe a repository-based system that we at The Data Base Approach have developed for physical database management. ■

DATABASE & CLIENT/SERVER WORLD

MARCH 26-28, 1996 • HYNES CONVENTION CENTER • BOSTON

FOR QUESTIONS

CALL: **508-470-3880**

24 HOUR FAX LINE: **508-470-0526**

EMAIL: **ConfReg@dciexpo.com**

When responding by EMail, please refer to KOADM01 and DBW in your subject line. Include your name, title, company, conventional mailing address, including mail stop, telephone and fax numbers.

☐ **YES!** Please send me a **FREE** Conference Brochure: (5488)

☐ Register me for the **FREE** 2-Day Exposition: (5489)

☐ Send me information on becoming an Exhibitor

MAIL THIS CARD TODAY!

Please add my colleague to your mailing list!

Name: Title:

Company:

Street: Mail Stop:

City: State: Zip:

Telephone: Fax:

EMail:

Name: Title:

Company:

Street: Mail Stop:

City: State: Zip:

Telephone: Fax:

EMail:

KOADM01

PLACE
STAMP
HERE

DATABASE & CLIENT/SERVER
WORLD

DCI
204 Andover St.
Andover, MA 01810

MARCH 26-28, 1996 • HYNES CONVENTION CENTER • BOSTON

DATABASE & CLIENT/SERVER WORLD

THE LARGEST AND MOST INFORMATIVE DATABASE & CLIENT/ SERVER EVENT OF THE YEAR!

In-Depth Coverage of Today's Hottest Topics:

- Data Warehousing & OLAP
- Parallel Databases & VLDBs
- Middleware
- Client/Server Tools and Application Development
- Systems Management for the Client/Server Environment
- Object-Oriented Technologies
- Multi-Tier Architectures
- Extending Relational DBMSs
- Client/Server Application Packages

and much more!

The ONLY Database & Client/Server Event Where You Can Hear From All The Top Experts Including:

- George Schussel
- Chris Date
- Jeff Tash
- Richard Finkelstein
- Phil Bernstein
- Dave McGoveran
- Herb Edelstein
- Ron Ross
- Shaku Atre

Special Presentations By:

Steve Mills
James Martin
John Donovan
Mitchell Kertzman
Jim Gray

For More Information Call (508) 470-3880 Today or Visit Our Home Page at http://www.DCIexpo.com/

Sponsored by:

DCI

KOADM02

PLUS

Client/Server Application Packages Conference

Presented by META Group

The Most Extensive Coverage Available Anywhere!

We disagree with the person who called our conference the "COMDEX of software development."



It's a lot cooler than COMDEX.

Register for Software Development '96
March 25–29, 1996, San Francisco, CA

It's cooler—and a lot more valuable for development professionals. So come rub elbows with everybody who's anybody in the field of software development.

Get up to speed on the latest tools for object-oriented programming, client/server migration and rapid application development. And take home skills you can use immediately.



With 200 classes and 300 tools vendors, SD '96 is the largest event of its kind. Covering everything from component to enterprise application development.

To learn more about registering for SD'96, call us at 800.441.8826 or 415.905.2702, or send us E-mail at sd96west@mfi.com. For instant information, visit our Home Page at http://www.mfi.com/sdconfs.

You'll be all set to attend the conference that features far more development tools than any other show. Now how cool is that?

More tools. More intelligence.

Software Development and the Software Development logo are trademarks of MFI. COMDEX is a trademark of Softbank COMDEX, Inc.

The Evolution of Client/Server CASE

HOW THE NEW CASE TOOLS PROMISE TO SIMPLIFY DATA MANAGEMENT AND AUTOMATE DEVELOPMENT.

Maurice Frank is DBMS*'s technical editor, based in Marietta, Georgia. You can reach him via CompuServe at 72167,736 or via the Internet at mfrank@mfi.com.*

By Maurice Frank

Years ago, pundits proclaimed the death of the CASE (Computer-Aided Software Engineering) market. Yet today, it's hard to conceive of a development project that does not use some kind of modeling tool to design, generate, and manage databases. Today's CASE tools are proving to be indispensable in managing complex data, application, and object designs. These tools are also assuming a greater role in the development and construction process. This article reviews CASE and modeling tools for client/server development. It also examines the emerging stratification of the CASE/modeling category into different classes of tools for managing different kinds of models: physical data models, logical data models, and object and process models.

How We Got Here

CASE tool support for client/server has not been an overnight phenomenon. It's taken several years for today's leading products to evolve to their present level of support. Looking back, it's easy to see several functional stages along the way. By definition, logical design is independent of the type of platform or architecture underlying the physical implementation. The most practical adaptations by CASE tools to client/server development have focused first on physical database design and more recently on application construction. Each side of the client/server equation has seen incremental advances in the degree of CASE tool integration.

CASE tools began their forays into client/server environments by generating, and later reverse-engineering, relational DBMS schema. At first, the tools translated entities and attributes into SQL CREATE TABLE statements and other Data Definition Language (DDL) statements stored in disk files. DBAs could then run the scripts on the DBMS server. As the tools matured, they employed native DBMS drivers to connect to a DBMS so they could run the DDL scripts and create DBMS tables and indexes. (The last two years have seen greater reliance on ODBC instead of native drivers.) As an encore following table creation, CASE tools such as ERwin from Logic Works began generating triggers and stored procedures, often from customizable templates. Along the way, several tools were designed to reverse-engineer existing database schemas into their CASE repositories, and some even lay out entity-relationship (ER) diagrams automatically after they catalog the relationships among tables.

But client/server systems do not live by databases alone. About two years ago, integrating CASE tools with application development tools became a hotbed of activity. Sybase's PowerBuilder uses extended attributes to capture interface properties of controls associated with table columns. Because PowerBuilder stores these attributes in relational tables, several CASE vendors enhanced their ER models to capture extended attributes. They could then export those extended attributes to PowerBuilder's extended attribute tables (many also import extended attributes into

a data model). This feature lets you specify interface properties during the modeling process. For example, you could specify a radio button control for a column. Several CASE vendors jumped on the PowerBuilder bandwagon. Gupta's SQL-Windows also attracted support from several CASE vendors.

Extended attributes are an important aspect of PowerBuilder development. CASE tools are only now beginning to generate PowerBuilder applications. One of the earliest examples is Sybase's own S-Designor, a CASE tool that Sybase acquired from SDP Technologies in 1995. Version 4.2.1 of S-Designor, which shipped in November 1995, generates DataWindows, windows, menus, and application objects. More important, developers are able to specify ancestors from which generated objects should inherit, including ancestors in third-party frameworks. S-Designor ships with templates that link Sybase's PowerBuilder Application Library to S-Designor. Sybase promises that future versions of S-Designor will support Visual Basic and other application development tools.

With their heads out of the PowerBuilder sands, CASE vendors are also eyeing other client/server development tools such as Visual Basic (especially the new Visual Basic 4.0 Enterprise Edition — see the cover story, "Visual Basic 4.0: Ready for the Enterprise?" on page 44), Borland's Delphi, and Forté Software's Forté. Both Logic Works and LBMS support Visual Basic 3.0, and both are now upgrading their tools to support Visual Basic 4.0. Other vendors that have already tied their tools to VB (or have immediate plans to do so) include Popkin Systems & Software, Sterling Software (which recently acquired KnowledgeWare and its ADW CASE toolset, which is now known as Key:Enterprise), Sybase, Select Software Tools, and Rational Software (the latter two vendors sell object-oriented CASE tools). I expect VB to attract the same broad support now enjoyed by PowerBuilder. Meanwhile, Borland developed a link between Delphi and LBMS's Systems Engineer, and Popkin Systems & Software has announced that it will support Delphi in an upcoming version of System Architect. Select Software Tools (Select Enterprise), Protosoft (Paradigm Plus), and Sterling have (or will soon) integrated their CASE tools with Forté.

Despite wider and deeper integration between CASE tools and client/server application development tools, current levels of integration are often little more than thin, snapshot-in-time bridges built between the two islands of design and construction. In theory, a design model is a high-level abstract view of a working implementation, yet today's tools rarely treat the model and the application or database as two views of the same thing. Today they remain two separate things that can exchange and translate relatively small parts of their world views. Typically, you create a design model and then generate or build the application in a one-way process.

Another force shaping today's CASE for client/server market is the greater participation by object-oriented (OO) CASE vendors. OO CASE tools embrace object analysis and design methods and notations, as opposed to ER and structured analysis and design diagrams (although some object methods and notations trace their lineage to earlier data and process models). OO CASE tools initially generated DDL for object-oriented DBMSs, but some also began mapping classes to RDBMS tables, reflecting the fact that many OO applications interact with an RDBMS, not an ODBMS. On the application side, the usual targets include C++ and Smalltalk, because these languages carry the flag in most OO camps. However, OO CASE vendors such as Rational and Protosoft have started to support PowerBuilder 4.*x* and Visual Basic 4.0, because each product steers developers toward greater reliance on classes as an integral part of its application structure. OO CASE vendors finally have a place to hang their hats in the 4GL world.

Several vendors have marketed very low-cost ($500 or less) versions of their CASE tools. These appeal to solo developers, and most support desktop products such as Microsoft Access or Xbase. Unlike their big brothers, who typically support team development, the desktop versions are usually single-user, and some features may also be limited. For example, Sybase's Star-Designor is a subset of S-Designor, but it supports only physical data models, not conceptual models. Logic Works also markets desktop versions of its ERwin product.

On the Horizon

What's next? At least three significant developments signal major changes in how CASE tools must support client/server DBMS and application development tools. On the server side, the next major challenge to client/server CASE vendors will be supporting new object features blended into RDBMSs such as the forthcoming Oracle8. You would expect that CASE tools with strong OO modeling features would be well-primed for this extension, but until DBMS vendors reveal more details regarding their object extensions, it is hard to say what will be needed and how much effort will be required to support it. I also expect object RDBMSs to have an impact on application development tools as well.

The modeling of decision-support and data warehouse applications is another frontier that few CASE vendors have explored, but I assume this market will open up within a year or two. Ralph Kimball, a *DBMS* contributing editor and data warehouse consultant, has discussed the inappropriateness of ER models for decision-support applications (see "Is ER Modeling Hazardous to DSS?" Ralph Kimball, *DBMS*, October 1995, page 17). Logic Works is investigating how to support (in a future version of ERwin) the modeling of star schemas, which are often used to represent data warehouse designs. Asymetrix is also weaving together models and queries with its new Info Assistant product. InfoAssistant is an end-user query tool that can create queries using information from an Asymetrix Info Modeler model.

On the application side, component thinking is slowly replacing systems thinking, especially as OLE becomes more common. CASE vendors must extend their tools to model and generate independent components in addition to integrated systems and applications. Again, I suspect that CASE tools supporting class modeling will have a head start in this arena. One example of this approach is Select Software Tools' Select Enterprise for Visual Basic. This object modeling CASE tool generates in-process and out-of-process OLE servers as part of a Visual Basic 4.0 application.

A housekeeping issue to watch is the future of open repositories. Today, several kinds of tools have their own storage mechanisms, although some tools interact with repositories used by other tools. Maintaining multiple representations of the same system is asking for trouble, so I would not be surprised to see more developers turn to repositories shared by multiple tools. LBMS recently introduced its Object Manager product, which serves as a repository for application objects created by tools such as PowerBuilder and Visual Basic. Object Manager also stores test scripts for SQA's TeamTest. I also expect vendors that sell DBMS servers, application development tools, and CASE tools to encourage developers to use RDBMSs as shared object and model repositories.

Categorizing CASE Tools

The client/server CASE tool market consists of tools that support the analysis, design, and modeling of client/server database systems. This market encompasses a diverse group of products with varying feature sets. Classifying tools along functional lines can help you choose an appropriate product for your environment. (Ignore vendor terms such as "high end" or "low end," because these rarely convey any useful or meaningful comparative information.) The following categories are based on broad feature sets:

- data modeling only
- data and process modeling
- object modeling

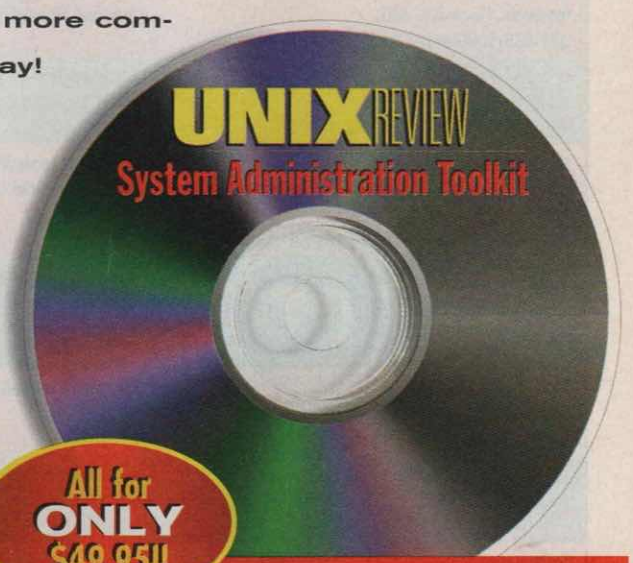
(continued on page 72)

Have the Tools of the Trade at Your Fingertips with *UNIX Review's*

System Administration Toolkit CD-ROM

The comprehensive, low-cost system administration solution!

To help lighten your workload, the *UNIX Review* technical staff has compiled on one CD-ROM the tools you need to perform most system and network administration tasks. These include security tools, languages, utilities, internet tools, e-mail tools, and performance monitoring tools. As an added bonus, we're throwing into the toolkit the code listings from all *UNIX Review* System and Network Administration supplements. You can't find a more complete set of tools for the work you do every day!



Comes Complete with the Following:

Security Tools
- kerberos
- tripwire
- cops
- securelive
- shadow
- Satan

Languages
- PERL 5.0
- Tcl/Tk
- expect

Utilities
- FAQs

Plus:
- Internet Tools
- E-mail Tools
- Performance Monitoring Tools
- And dozens more!

ADDED BONUSES:
- Code listings from *UNIX Review's* System and Network Administration supplements
- System administration product listings from *UNIX Review's 1995 Buyer's Guide*

All for ONLY $49.95!!

ORDER FORM

YES! Rush me ____ copies of the *UNIX Review* System Administration Toolkit CD-ROM at $49.95 each

Quantity	Price	Total
	$49.95	$
	Subtotal	
	Shipping/Handling	
	Sales Tax	
	TOTAL	**$**

ORDER TODAY!

Sales Tax: NY residents-8%, IL residents-6.25%, CA residents-7.5%, TX residents-7.5%, GA residents-5%
Shipping & Handling: $4 for U.S. and Canada orders; $10 for international orders

Complete this form and return to: *UNIX Review* System Administration Toolkit CD-ROM, 1601 West 23rd Street, Suite 200, Lawrence, KS 66046. Telephone: 913/841-1631; FAX: 913/841-2624; e-mail: sysadmncd@mfi.com

Name _____

Address _____

City _____ State _____ Zip _____

Phone _____

❑ Check ❑ VISA ❑ Mastercard

Credit Card #: _____ Exp. Date _____

Signature: _____

TABLE 1. Product Chart

PRODUCT/COMPANY	DESCRIPTION
Bachman Information Systems Inc., Burlington, Mass. 617-273-9003 or fax 617-229-9904	Several tools, including GroundWorks for data and process modeling, Terrain for database generation, Ellipse for application development, and WindTunnel for database performance monitoring
CAST Workbench Cast Software Inc., San Francisco, Calif. 415-284-7970 or fax 415-284-7979	A collection of programs, including CW-Designer, an entity-relationship modeling tool, and several database administration programs
Chen Workbench Chen & Associates Inc., Baton Rouge, La. 504-928-5765 or fax 504-928-9371	A suite of tools for entity relationship diagramming, data flow diagramming, normalization, schema generation, and reverse engineering
Composer, Arranger Texas Instruments Software, Plano, Texas 214-995-2011 or fax 214-575-4144	An integrated set of modeling and application development tools
Navigator Systems Series Ernst & Young LLP, Dallas, Texas 214-665-5000 or fax 214-665-5555	A methodology-based application development environment that supports application packages such as SAP R/3
Designer/2000 Oracle Corp., Redwood Shores, Calif. 415-506-7000 or fax 415-506-7200	Data and process modeling tools that integrate with Oracle's Developer/2000 application development tools
EasyCASE Professional Evergreen Software Tools Inc., Redmond, Wash. 206-881-5149 or fax 206-883-7676	Integrated data and process modeling (several diagram types), schema generation, reverse engineering, and support for PowerBuilder
ERwin/ERX Logic Works Inc., Princeton, N.J. 609-252-1177 or fax 609-252-1175	A data modeling tool with versions that integrate with PowerBuilder, Visual Basic, and SQLWindows
Excellerator II Intersolv, Rockville, Md. 301-838-5000 or fax 301-838-5432	A data, process, and object modeling CASE tool that integrates with Intersolv's APS development tool and PowerBuilder
InfoModeler, InfoAssistant Asymetrix Corp., Bellevue, Wash. 206-462-0501 or fax 206-637-1504	InfoModeler is a data modeling tool based on Object Role Modeling; InfoAssistant is an end-user query tool that can build queries based on an InfoModeler data model
Key:Enterprise Sterling Software, Applications Development Division, Atlanta, Ga. 800-444-8575 or fax 770-364-9454	A suite of tools including data and process modeling (Key:Model) and application development (Key:Assemble). Also integrates with PowerBuilder and Visual Basic
Systems Engineer, Process Engineer, Object Manager LBMS Inc., Houston, Texas 713-625-9300 or fax 713-625-9480	Systems Engineer is a data and process modeling tool that integrates with PowerBuilder, Visual Basic, and other tools; Process Engineer is a methodology and project management tool; Object manager is a shared repository
Paradigm Plus Protosoft Inc., Houston, Texas 713-480-3233 or fax 713-480-6606	An object modeling CASE tool that generates PowerBuilder and Forté applications
Rational Rose Rational Software Corp., Santa Clara, Calif. 408-496-3600 or fax 408-496-3636	An object modeling CASE tool with versions that generate PowerBuilder and SQLWindows applications; an upcoming version will generate Visual Basic applications
S-Designor Sybase Inc., Concord, Mass. 508-287-1500 or fax 508-369-4695	A data modeling tool that generates PowerBuilder application objects (DataWindows, menus, windows); generated objects can inherit from objects in application frameworks (future versions will support Visual Basic and other development tools)
Select Enterprise Select Software Tools Inc., Santa Ana, Calif. 714-957-6633 or fax 714-957-6219	An object modeling CASE tool with versions that generate applications for Forté and Visual Basic
Silverrun-Enterprise Computer Systems Advisers Inc., Woodcliff Lake, N.J. 201-391-6500 or fax 201-391-2210	Four integrated tools for data and process modeling; generates applications for PowerBuilder, SQLWindows, Blyth Software's Omnis 7, and others; runs on several platforms
System Architect Popkin Systems & Software Inc., New York, N.Y. 212-571-3434 or fax 212-571-3436	Supports several diagram types and notation styles for data, process, and object modeling; also integrates with PowerBuilder and SQLWindows; support for Borland's Delphi has also been announced
Visible Analyst Workbench Visible Systems Corp., Waltham, Mass. 617-890-2273 or fax 617-890-8909	A data, process, and object modeling CASE tool. Integrates with PowerBuilder and SQLWindows

A representative sample of CASE tools that support client/server database and applications development.

Learn how to develop and maintain your Very Large Database at:

The 2nd Annual VLDB Summit

The Second Annual Conference Devoted Exclusively to Very Large Commercial Databases

April 29 - May 1, 1996
Chicago, IL

EXPERT SESSIONS ON:

- VLDB Performance Design
- Large Decision Support Case Studies
- Techniques for Huge Data Warehouses
- High Availability and Recovery
- Parallel Databases: SMP, MPP
- Data Mining and Business Intelligence
- VLDB Upsizing Tips for Oracle, Informix, Sybase and DB2

Sponsored by:

DATABASE PROGRAMMING & DESIGN

DBMS DATABASE & CLIENT/SERVER SOLUTIONS

un Miller Freeman A United News & Media company

Phone: (415) 905-2267
or Mail to:
VLDB
600 Harrison Street
San Francisco, CA 94107

For information FAX to: (415) 905-2218

❑ Yes! I'm interested. Please send the 2nd Annual VLDB Summit catalog when available.

❑ Yes! I'll be there. Please register me for the 2nd Annual VLDB Summit in Chicago, IL on April 29 - May 1, 1996.

Name

Title

Company

Address

City State Zip

Phone Fax

DBS95

(continued from page 68)

- methodology-based tools
- other

Data Modeling Tools

Some of the most popular client/server CASE tools limit their scope to data modeling. Two examples are the ERwin product family and S-Designor. Both ERwin and S-Designor use ER diagrams to capture database designs, and both offer strong support for schema generation, reverse engineering, and integration with front-end application development tools such as PowerBuilder. However, a key difference is that S-Designor lets analysts model separate logical and physical designs, while ERwin supports only one kind of model at a time. You can use ERwin to design either a logical model or a physical model, but you would have to synchronize them manually. S-Designer can generate a physical model from a conceptual model, and it will synchronize the two while preserving developer decisions such as denormalized attributes. As an ERwin user, I have found that its strong support for physical database design sometimes makes it easy to drift away from a logical focus in favor of server-specific requirements.

Most data modeling CASE tools employ various flavors of ER diagrams. An exception is Asymetrix's InfoModeler (see "Black Belt Design," *DBMS* Interview, September 1995, page 38), which uses Object Role Modeling. InfoModeler graphically represents a design as objects that play roles, and it can also express a design as a collection of English-like fact statements. Many developers have found this feature useful for verifying designs with non-technical end users.

Although several CASE tools generate schemas for RDBMSs, Cast Software's CAST Workbench is one of the few to integrate data modeling with other database administration tools. CW-Designer, an ER modeling tool, is one of six applications in CAST Workbench. DBAs use the other modules to reverse-engineer, write, test, and manage scripts, and to migrate data into Sybase SQL Server and Microsoft SQL Server databases. Because many client/server projects are still performed by small teams, the same developers often handle a broad array of tasks ranging from database administration to systems design.

Process Modeling Tools

Data shouldn't just sit around. It must do *something*. Consequently, client/server CASE tools help developers design programs to manage data. Traditional process modeling techniques such as data flow diagrams, structure charts, and state transition diagrams may seem like dinosaurs in our object-oriented age. But, unlike a Raptor, these techniques are still alive and kicking in many CASE tools. System Architect supports more than a dozen different process modeling diagram styles, along with several flavors of ER diagrams and several different object modeling diagrams. Other products that integrate process and data modeling include Silverrun-Enterprise from Computer Systems Advisers and EasyCASE Professional from Evergreen Software Tools.

Blending data, process, and object diagrams with a common repository lets developers design systems from different viewpoints using various levels of abstraction. For example, LBMS integrates several products around its repository. Systems Engineer is a modeling tool that also generates database schemas (from the SE/ServerBuilder module) and application components. It also models GUI applications using diagrams that illustrate navigation paths among windows or forms. Process Engineer provides online methodology guides as well as project management features. Also, Object Manager extends LBMS's repository technology to application development tools, testing tools, and other development products, so they can benefit from repository-based documentation, versioning, check-in, and check-out controls.

ADVERTISER DIRECTORY

RS#	Company	Page
	Apple Computer	15
55	Aspect Software Engineering	28
78	Atre Associates, Inc.	52
119	Automation Technology	81
	Blenheim NDN	38
135	Borland International	C4
278	Bradmark Technologies	21
	Business Rules Conference	75
	Cast Software	78
188	Computer Associates	7
	Crystal Services	51
62	Dharma Systems	77
	Digital Consulting, Inc.	65
251	Embarcadero Technologies	80
117	EmbarcaderoTechnologies	82
123	Inquiry.com	23
123	Inquiry.com	25
123	Inquiry.com	27
86	Intersolv	4–5
116	IQ Software	10–11
139	JYACC	58
84	Knowledge Based Systems	47
200	Logic Works	29
40	Magic Software Enterprises	49
155	Mdbs	19
202	Microrim	54
	Microsoft	C2–1
	Microsoft	37
	Microsoft	43
68	MITI	C3
	Oracle Corporation	9
57	Platinum Technology	13
118	Popkin Software	57
227	Powersoft Corporation	35
115	Powersoft Design Products Div.	24
15	Protoview	22
159	Quadbase Systems, Inc	36
38	Raima Corporation	33
56	Red Brick Systems	17
26	Rogue Wave	53
9	Sequiter Software	79
	Simba Technology	39
	Software Development Conf.	66
94	SPFC	77
76	SQA	41
102	Subtle Software	77
73	Sylvain Faust Inc. - SFI	83
5	Syware	77
28	Trinet	77
66	Visigenic	2
	VLDB Conference	71
121	Working Set	40
77	Wrox Press	77
61	Zitel Corporation	30

DBMS SALES CONTACTS

Publisher
Philip Chapnick (415) 655-4181

East Coast Regional Representative
Len Ganz (617) 244-8240

Account Representative, Eastern Region
Paul Calento (617) 244-8250

West Coast Regional Representative
Susan Dawson (415) 655-4242

Account Representative, Western Region
Joseph DeSantis (415) 655-4309

Advertising Coordinator
Gae Reck (415) 655-4245

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

READER SERVICE CARD • FREE INFORMATION

1. Fill in your name and address on this card.
2. Circle the numbers corresponding to the advertisement, article, or directory listing that interests you. Use the Advertiser Directory to the left as an easy reference.
3. Just drop the card in the mail. Postage is paid by *DBMS*. Your requested information will be mailed to you free of charge by the vendor.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone (_____)_____

FAX (_____)_____

This January issue card void after April 30, 1996

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108
109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126
127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144
145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162
163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180
181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198
199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216
217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234
235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252
253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270
271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288
289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306
307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324
325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342
343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360
361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378
379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

READER SERVICE CARD • FREE INFORMATION

1. Fill in your name and address on this card.
2. Circle the numbers corresponding to the advertisement, article, or directory listing that interests you. Use the Advertiser Directory to the left as an easy reference.
3. Just drop the card in the mail. Postage is paid by *DBMS*. Your requested information will be mailed to you free of charge by the vendor.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone (_____)_____

FAX (_____)_____

This January issue card void after April 30, 1996

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108
109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126
127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144
145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162
163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180
181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198
199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216
217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234
235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252
253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270
271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288
289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306
307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324
325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342
343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360
361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378
379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

READER SERVICE CARD • FREE INFORMATION

1. Fill in your name and address on this card.
2. Circle the numbers corresponding to the advertisement, article, or directory listing that interests you. Use the Advertiser Directory to the left as an easy reference.
3. Just drop the card in the mail. Postage is paid by *DBMS*. Your requested information will be mailed to you free of charge by the vendor.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone (_____)_____

FAX (_____)_____

This January issue card void after April 30, 1996

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108
109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126
127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144
145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162
163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180
181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198
199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216
217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234
235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252
253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270
271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288
289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306
307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324
325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342
343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360
361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378
379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO. 200 PITTSFIELD, MA

POSTAGE WILL BE PAID BY ADDRESSEE

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

P.O. Box 5282
Pittsfield, MA 01203-9845

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO. 200 PITTSFIELD, MA

POSTAGE WILL BE PAID BY ADDRESSEE

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

P.O. Box 5282
Pittsfield, MA 01203-9845

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO. 200 PITTSFIELD, MA

POSTAGE WILL BE PAID BY ADDRESSEE

DBMS
DATABASE & CLIENT/SERVER SOLUTIONS

P.O. Box 5282
Pittsfield, MA 01203-9845

If specialized titles weren't so effective, we wouldn't be the world's fastest growing high-tech publisher.

But we are.

It's funny how things can sneak up on you. Like Miller Freeman's strong growth in high-tech magazines.

As the high-tech world has become increasingly specialized, Miller Freeman has quietly assembled an outstanding selection of titles to fit this reality. Since September 1992, we have added *Cadence, Dr. Dobb's Journal, Microsoft Systems Journal, DBMS and OS/2 Developer. Software Development, OS/2 Magazine, Stacks: the Network Journal* and *Game Developer* were launched in that same time frame.

We work hard to identify important markets and then deliver magazines with tight editorial focus and highly targeted audiences. This gives you access to in-depth articles and product reviews—coverage you won't find in broader magazines.

Our specialized expertise has grown to include international marketing and trade shows. Miller Freeman is one of North America's largest producers of high-tech events featuring SD East & West, OS/2 World, Autodesk University and the Computer Security Conference.

So next time you're looking for targeted high-tech information, look for a specialized magazine from Miller Freeman.

mf Miller Freeman, Inc.
A MEMBER OF THE UNITED NEWSPAPERS GROUP



Recently, CASE tools have shifted the focus of their process modeling techniques along two lines: object-orientation and business process reengineering. In addition to Popkin's System Architect, Intersolv's Excellerator II and Visible Systems' Visible Analyst Workbench (VAW) supplement process diagrams based on structured techniques with new diagram types for modeling classes and objects.

Object Modeling Tools

Object orientation is not simply a different way to program, it's also a different way to analyze, design, and model systems. While numerous object modeling approaches have emerged over the last five years, about a half dozen have pulled ahead of the pack to gain widespread marketshare. Following in their footsteps, a growing collection of object modeling CASE tools support one or more object modeling diagram and notation styles. These techniques are often named after their originators, including Grady Booch, James Rumbaugh, Ivar Jacobson, Peter Coad and Ed Yourdon, and others.

As today's application development tools become increasingly object-oriented, object modeling CASE tools become more relevant to application developers who do not use C++ or Smalltalk, the traditional focus for object CASE tools. Some popular object CASE tools include Rational Rose by Rational Software, Paradigm Plus by Protosoft, and Select Enterprise by Select Software Tools, which is based on Select OMT Professional. All three now generate application objects for development tools such as PowerBuilder, Gupta SQLWindows, Forté, and Visual Basic.

Methods and Manpower

In some ways, a CASE tool is like a word processor: If you don't know how to write, not even the richest feature set will turn you into a Shakespeare. Even the best CASE tool cannot turn you into a great systems designer if you don't understand the basics of analysis and design. Some CASE vendors, such as LBMS, provide methodology modules that guide developers through the recommended steps in a development project. (LBMS recently acquired Corporate Computing Inc., which markets a methodology geared to developing GUI applications.) The methodologies are usually customizable, so companies can tailor them where necessary.

Beyond online methodology aids, several consulting firms, such as Arthur Andersen and Ernst & Young, provide onsite counselors with the CASE tools they sell. These mentors advise clients not only on the mechanics of making the CASE tool work, but on the fundamental procedures and techniques for developing large systems that involve teams or leagues of developers. Many of the large shops using these services began using CASE tools on mainframe projects, and now utilize them on client/server systems as well. Among the larger shops, Texas Instruments' IEF CASE tool is a popular choice for enterprise-wide development. IEF, supplemented by Texas Instruments' Composer object development environment and Arranger object assembly environment, provides a full life cycle development solution with an implicit methodology.

Building Bridges

Not every tool fits into a classification scheme, and sometimes the most interesting ones are the hardest to categorize. One example is LEAD (Logical Entity Action Design), by Client Server Designs (415-321-5338), a Menlo Park, Calif.-based consulting and applications development firm. LEAD is a bridge between ERwin or System Architect and SQLWindows. LEAD generates a working SQLWindows application, the database, and test data.

Developers build a LEAD model using System Architect or ERwin. The model includes ER diagrams, and if System Architect is used, developers can also create entity state transition diagrams. LEAD's linchpin is the use of two-character prefixes for column names in ER diagrams. These prefixes signify a domain such as date, user ID, state, phone, and so on. LEAD provides almost 20 built-in domain prefixes, and developers can add new ones.

The LEAD Toolkit Console program reads a text file exported by the CASE tool. LEAD generates an application using the LEAD class library, a set of SQLWindows QuickObjects. Each entity in the model becomes a table in the database and a form in the application. (LEAD can also generate one-to-many forms.) Each field on a form is a SQLWindows QuickObject that is based on a LEAD domain.

Client Server Designs is currently working on LEAD version 2, which will generate complete applications for Visual Basic, PowerBuilder, and C++. It will also create OLE controls, so any development environment that supports OLE controls will be able to benefit from LEAD's ability to translate model specifications into application components. Client Server Designs plans to ship version 2 in late 1996.

CenterView Software's (South San Francisco, Calif., 415-873-1295) Choreo is another example of a third-party tool that builds applications from database schemas. Choreo is not a CASE tool, but it performs functions that some CASE tools are now adopting. Choreo reads a DBMS catalog and imports information about tables, columns, primary and foreign keys, indexes, and relationships in a database. If the database is well-designed, Choreo can infer most of the relationships among tables, and it will prompt the developer when it cannot. Choreo stores this information in a "model" file, which is more like a database of its own because no graphical diagram is involved.

A developer runs the Choreo DataLink Manager along with Visual Basic, and drags and drops fields from the Choreo model onto controls on a Visual Basic form. Choreo includes its own grid control, which responds to the drag and drop of a field from the Choreo model. Developers can use other third-party grid controls, but they must write calls to the Choreo API. Once Choreo learns the linkage between a field and a control, it manages all interaction between the Visual Basic application and the database. This approach enables Choreo to coordinate data navigation across one or many windows without additional programming. Choreo includes its native database drivers for Oracle, Sybase and Microsoft SQL Server, and Gupta's SQLBase. An earlier version of Choreo for Gupta's SQLWindows is now marketed by Gupta.

Logic Works is developing a version of ERwin for Visual Basic that will create the Choreo model file from an ERwin model. ERwin will also generate Visual Basic 4.0 forms that use the Choreo controls and database engine.

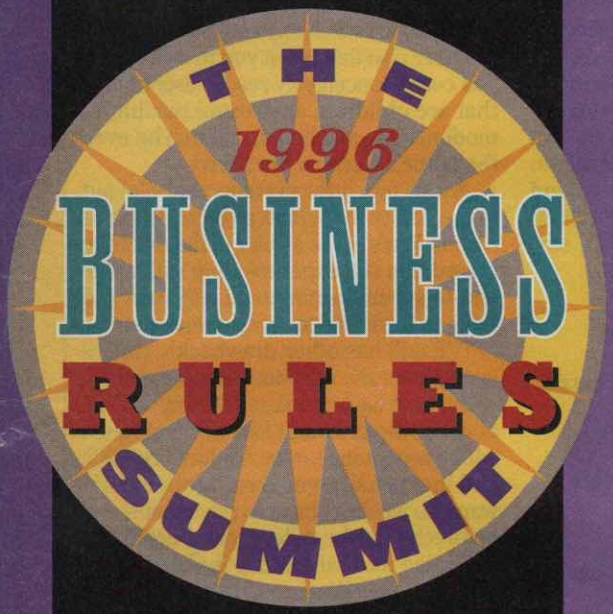
Designing Complexity

Although CASE tools capture the extensive detail necessary to generate the server side of a client/server database, you'll still need a good DBA tool for many administration tasks. Now, CASE tool vendors are capturing more information that they use to generate 4GL application components for several popular front-end development tools. This makes client/server CASE tools more valuable because they enable developers to spend more time and effort thinking about what a system should do, instead of laboring over the mechanical construction of forms, menus, and related objects. After all, doing the right thing is as important, if not more important, than perfecting or optimizing its implementation.

Client/server systems are increasingly complex. Even if you use rapid prototyping to gather user feedback regarding interface design and application function, any non-trivial system needs a well thought-out design. CASE tools enable developers to model database applications and integrate with both back-end DBMS servers and the front-end application development tools. If you're developing a sophisticated client/server system without using a CASE tool, you may be doomed for extinction. (For an extensive list of CASE tools, see Table 1, page 70.) ■

Capturing the Evolving Business Process

San Francisco, CA
February 26-28, 1996



Spurred by advances in technology, businesses are moving faster than ever to exploit market opportunities. As a modeler, designer, or developer of information systems, you must meet this challenge or risk letting your systems fall behind.

The most effective way for you to learn how to build systems that accurately represent and support a rapidly changing business is by attending THE BUSINESS RULES SUMMIT. *Find out how experienced practitioners are using business rules to define key business processes; build enterprise models and systems; and add meaning to the organization's vital self knowledge. Using business rules, you can forge a more productive relationship between IT and business users.*

Attend These Expert Sessions and more !

Plus
in-depth sessions from a selection of vendors active with business rules!

- **"An Introduction to Business Rules," Ronald Ross, Ross & Associates**
- **"Business Rules and Object Design," Barbara von Halle, Spectrum Technology Group**
- **"Achieving Enterprise Cognition Through Intelligent Objects," David Taylor, Enterprise Engines**
- **"Beyond Reengineering: How Businesses Learn to Evolve," Michael Rothschild, The Bionomics Institute**
- **"Understanding Business Objects," Robert Shelton, Open Engineering Inc.**
- **"Business Rules: Building the Semantic Network Database," David Britton, Britton Lee Inc.**
- **"When Rules Become Stored Procedures: Managing Database Server Objects," Craig Mullins, Platinum Technology Inc.**
- **"Data Modeling By the Rules...Or Was That 'Next to' the Rules?" David Hay, Group R Inc.**
- **"Digging for Business Rules in Legacy Systems," William Ulrich, Tactical Strategy Group**
- **"The Conscious Database: Implementing Neural Networks with Stored Procedures," Michael Gonzalez, The Focus Group Ltd.**
- **"Practical Implementation Techniques for Business Rules in Three-Tier Architectures," Tim Quinlan, Quinlan Consulting**
- **"Case Study: Business Rules in Production," Sandra Geyer, Metropolitan Water District of Southern California**

YES, PLEASE SEND FREE INFORMATION ON ATTENDING!

For information fax to 415-905-2218; mail to: The Business Rules Summit, 600 Harrison St., San Francisco, CA 94107 or call 415-905-2267. Internet: dstodder@mfi.com.

Name _____ Title _____

Company _____

Address _____ City _____ State _____ Zip _____

Phone _____ Fax _____ E-mail _____

DBS95

Sponsored by: DATABASE PROGRAMMING & DESIGN | DBMS | Miller Freeman, A United News & Media company

LETTERS

(continued from page 6)

So is his suggestion that we wind up with more types than we can handle; the number of types needed in his example having to do with lengths is exactly *two*.

- His description of how SQL "domains" behave is plain wrong.
- There *is* no logical difference between "an enumerated set" and "a set defined by a rule."

To close on a positive note: Celko's letter does include one point on which (I think) we are partly in agreement! He suggests that it might be possible to express a given calculation in two distinct but equivalent ways, one involving an explicit multiplication of two weights and the other not. So if we've "disallowed the operation weight times weight at the data type level," one of these two distinct but equivalent expressions might be legal and the other not.

I agree this is a valid concern (at least in general, though I'm not sure it's valid in the case of multiplying two weights in particular). I've pointed out elsewhere that, for example, if we can't add two dates together, we obviously can't go on to divide that sum by two to obtain the date midway between the given ones — yet this latter would be a useful and valid thing to do. (This criticism applies to SQL, incidentally.) But the moral isn't that types should simply "inherit" operations from their representations, as Celko suggests, but rather that data type definition needs to be done carefully. What's more, this moral applies at least as much to built-in types as it does to user-defined ones.

C.J. Date
Healdsburg, Calif. ■

IVR *(continued from page 56)*

kit, and FoxExpress, Neon Software's data dictionary-driven application generator.

I wrote the system for a manager of a department full of consultants who work off-site. The consultants are very good at what they do, but are practically brain-dead when it comes to time reporting. The department has just instituted a policy that requires employees to report time spent on projects at the end of each day. Employees call in from anywhere in the country to update the Time and Attendance database.

You can view sample code on the *DBMS* Web site (http://www.dbmsmag.com/IVR.html). The code is also available in our CompuServe forum library 13. This sample shows code fragments from the system related to IVR. (I left out processing details and error-trapping code to keep the discussion focused on VTK functions.) I started out by requesting the department secretary, who has a professional-quality voice, to record the system's various prompt files. The PC I used did not have a sound card, so, instead of VoysSmith, I used the telephone and the VoysAccess Recorder program (written in FoxPro) to record the prompts.

The sample code is heavily commented, so you should not have any problem following it. I do want to point out two superb facilities in VoysAccess and FoxPro that make IVR application development with the VTK relatively painless: the VSpeak( ) function and FoxPro's memo field facility. The toolkit's VSpeak( ) function is number-smart; that is, it can speak numbers properly in context. It can speak numbers normally (123 is "one hundred twenty-three"), numbers as digits (123 is "one, two, three"), numbers as days of the week (1 is "Sunday"), or numbers as money (12.34 is "twelve dollars and thirty-four cents"). To take advantage of this facility, you precede the number with the appropriate number "speak mode," such as DIGITS, DAY, or DOLLARS.

In many other IVR "languages," you would have a disk cluttered with voice files if the application allows callers to leave voice messages. In FoxPro, you can store voice messages in memo fields instead of individual disk files; FoxPro memo fields can store binary data as well as text. If the application supports multiple lines and runs on a network, you can avoid file-contention problems by not recording directly into the table's memo field. In the sample code, I dance around the problem by recording into a cursor (a temporary table, for all you brain surgeons who are not into databases).

■ A nice touch for systems that dial out is the ability to include a voice file of the respondent's name as part of the database record.

You should consider carefully whether there is an overriding business benefit to allowing callers to leave messages. Messages require more attention and they can use a lot of space, so you should impose a maximum recording time. Incidentally, the VTK uses voice files with the .vox extension. These voice files are not the same as .wav files and are more space-efficient.

I should mention that Voysys now also has a Visual Basic version of the VTK that comes bundled with a visual rapid application development tool. Visual Basic is the language of choice for many PC-based IVR development packages.

Bullet-Proofing Your Application

I have always preached that a system must anticipate almost everything that can go wrong to enable it to remain in control. IVR systems that dial out require the most attention. You, too, may discover that the application works correctly when connecting to home telephones, but fails to detect phone pickups or hangups when calling PBX (public branch exchange) phones. Remember also that when you're calling out, you could conceivably reach respondents that are children, answering machines, modems, fax machines, and maybe even the home's German Shepherd.

A nice touch for systems that dial out is the ability to include a voice file of the respondent's name as part of the database record. I'm still experimenting with some text-to-speech commercial offerings and hope to have one to recommend soon.

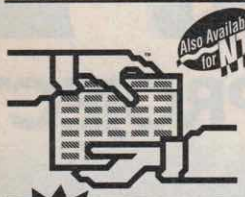
IVR still has other drawbacks. Unless we reinvent the telephone keypad, people will always have problems dealing with alphabetic input. I think IVR will gain much wider acceptance once we have reliable and inexpensive voice recognition software, so we don't have to rely exclusively on touch-tone for caller input. People like to remind me that this day is already here and offer the telephone company's directory assistance service as proof.

In my hometown of Washington, D.C., for example, you can speak the name of a Washington metropolitan city, and the directory assistance IVR application recognizes it. But you know what? That's not high-tech; au contraire, it is the ultimate in low-tech. There is a real person listening to that conversation and punching the correct button. So why doesn't the person just answer the phone and give out the requested number? Because the telephone company learned very quickly that every city has its share of lonely hearts just dying to talk to a real person, even if it happens to be a directory assistance operator. How boring! ■

- Neon Software, 335 N. Superior, Toledo, OH 43604; 419-255-6366 or fax 419-255-6371.
- Periphonics Corp., 4000 Veterans Memorial Highway, Bohemia, NY 11716; 516-467-0500 or fax 516-981-2689.
- Voysys Corp., 48634 Milmont Dr., Fremont, CA 94538; 510-252-1100 or fax 510-252-1101.

Dr. DeeBee™ ODBC Tools

Tools for ODBC development and support



Ever wonder why your ODBC app is not working? Why it's just too slow? If the ODBC driver is OK? Dr. DeeBee utilities reveal the inner workings of ODBC.

Call for a *Free* ODBC prescription!

NEW **Connect proprietary databases to Access, Visual Basic, and PowerBuilder with our Dr. DeeBee ODBC Driver Kit**

SYWARE Inc.
P.O. Box 91 Kendall, Cambridge, MA 02142
617-497-1376
Fax 617-497-8729
http://www.syware.com/drdeebee/

CIRCLE READER SERVICE NUMBER 5



Wireless DBMS Solution At Your Fingertips

WiceNet is a mobile and wireless service bureau providing paging, fax, electronic mail and database access services over wireless digital packet and cellular digital packet networks. **WiceNet** is based on the newest client-agent-server architecture to provide the most efficient way in wireless data communication and messaging.

Call for a hands-on free demo program of wireless interactive database query, mobile field sales application, and our wireless Internet e-mail connection service, **WiceMail**, today. See the performance and the usability of the new digital wireless data communication technology without spending thousands of dollars for system set-up.



5718 Westhiemer, Suite 1314
Houston, Texas 77057 USA
Tel: 713-785-9962
Fax: 713-785-9964

TriNet COMMUNICATIONS, INC.

CIRCLE READER SERVICE NUMBER 28

SPFC

Box 163
Madison, NJ 07940
800-676-7732
201-765-0791

SQL Accounting For Windows

- *SQL Client/Server Database*
- *Graphical User Interface*
- *Full Feature Accounting Package*
- *Written in Gupta Corp's SQLWindows*
- *Trial Copy — $100 per 3 modules for 30 days*
- *$2000 PER MODULE FOR RUN TIME*
- *Source Available*

* *General Ledger*
* *Accounts Payable*
* *Accounts Receivable*
* *OE Billing*
* *Inventory*
* *Purchasing*
* *Payroll*
* *Manufacturing*

CIRCLE READER SERVICE NUMBER 94

LEARN HOW TO DEVELOP GREAT POWERBUILDER APPLICATIONS

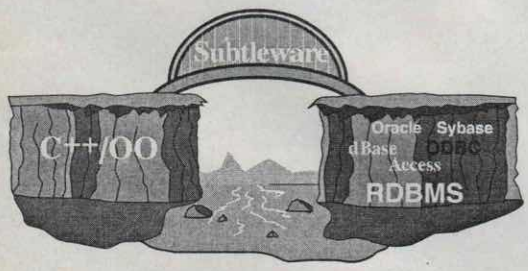


* Fully comprehensive guide to Powersoft's Client/Server package
* Full coverage of DDE, OLE, ODBC, Object Oriented Powerscript
* Learn how to use embedded SQL to query your databases
* Develops a fully functional commercial application
* Includes detailed advice on Sybase connectivity
* CD-ROM contains source code and applications

Author: **Prasad Bodepudi** ISBN 1-874416-60-5 **$49.95**
For more information or a free catalog of Wrox Press programming titles call **800 814 4527** or email **rachelm@ix.netcom.com**
WORLD CLASS PROGRAMMING BOOKS

WROX
TO ORDER
Tel. 800 937 5557
Fax. 800 PRI ORDER
QUOTE CODE **CL38** FOR FREE FREIGHT VISA MASTERCARD AMEX DISCOVER

CIRCLE READER SERVICE NUMBER 77



C++ Developers

Jump start your application development with the

Subtleware® C++ Class Generator

- Saves 30% of the time required to write C++ database applications
- Coverts RDBMS tables into C++ classes with the click of a button
- Runs under Windows 3.1, Windows NT, and Windows 95
- Supports Microsoft, Rogue Wave, Borland and other class libraries
- Create C++ objects that read and write themselves to your relational database
- Use an intuitive GUI that can be mastered in minutes
- Work with all popular databases including Access, Oracle, SQL Server
- Read/Write objects using an included DB class library, DBTools.h++, or MFC

Simply connect to a database, select schema in the database, and generate your C++ database classes!

Designer Edition **$129**
Developer Edition **$399**



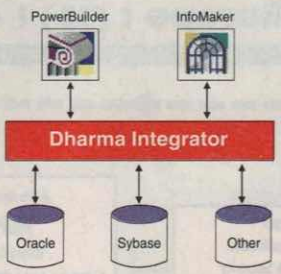
Subtle Software, Inc
7 Wells Ave.
Newton, MA 02159
P: (617) 558-4100
F: (617) 558-4103
email: subtle@world.std.com
http://world.std.com/~subtle/index.html

CIRCLE READER SERVICE NUMBER 102

Powering Tools With Data Integration

Dharma Integrator™ combined with PowerBuilder and InfoMaker, allows users to query, and create reports that join data from multiple and disparate data sources.

In the real world, data doesn't reside in one convenient location. Critical data exists throughout the enterprise like pieces in a puzzle. With Dharma Integrator, experience the power of integrating data scattered throughout your organization: data important to your company's needs, from accounting, marketing, sales, and manufacturing.



Find out how Dharma Integrator can solve your data integration needs. **Call Dharma today at 1-603-886-1400.**

Dharma Systems, Inc.
15 Trafalgar Square, Nashua, NH 03063

dharma

Dharma Integrator is a trademark of Dharma Systems Inc. Other trademarks are the property of their respective owners.
© 1995 Dharma Systems Inc.

CIRCLE READER SERVICE NUMBER 62

BREATHE LIFE INTO YOUR SYBASE™ PROJECT

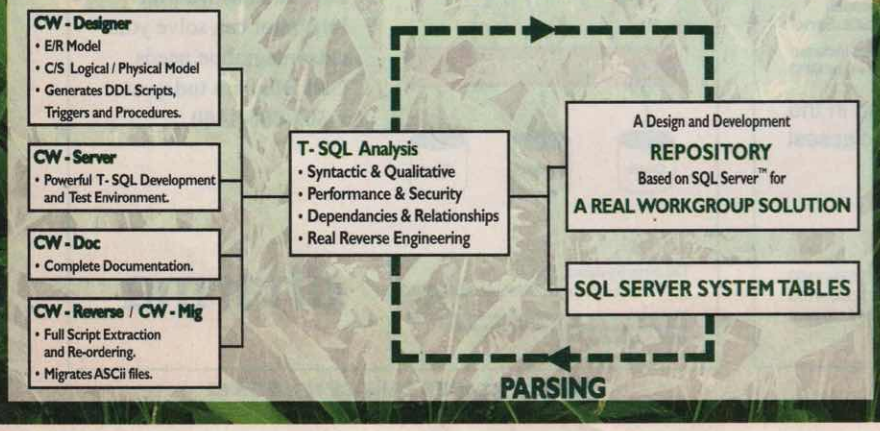
With CAST WORKBENCH™

THE BEST MULTI USER DESIGN & DEVELOPMENT ENVIRONMENT FOR SQL SERVER™

- Takes your Project from Design to Production.
- Guarantees Perfect Cohesion between all Phases.

CALL NOW !
for a Free Evaluation Kit
or a Technical White Paper.

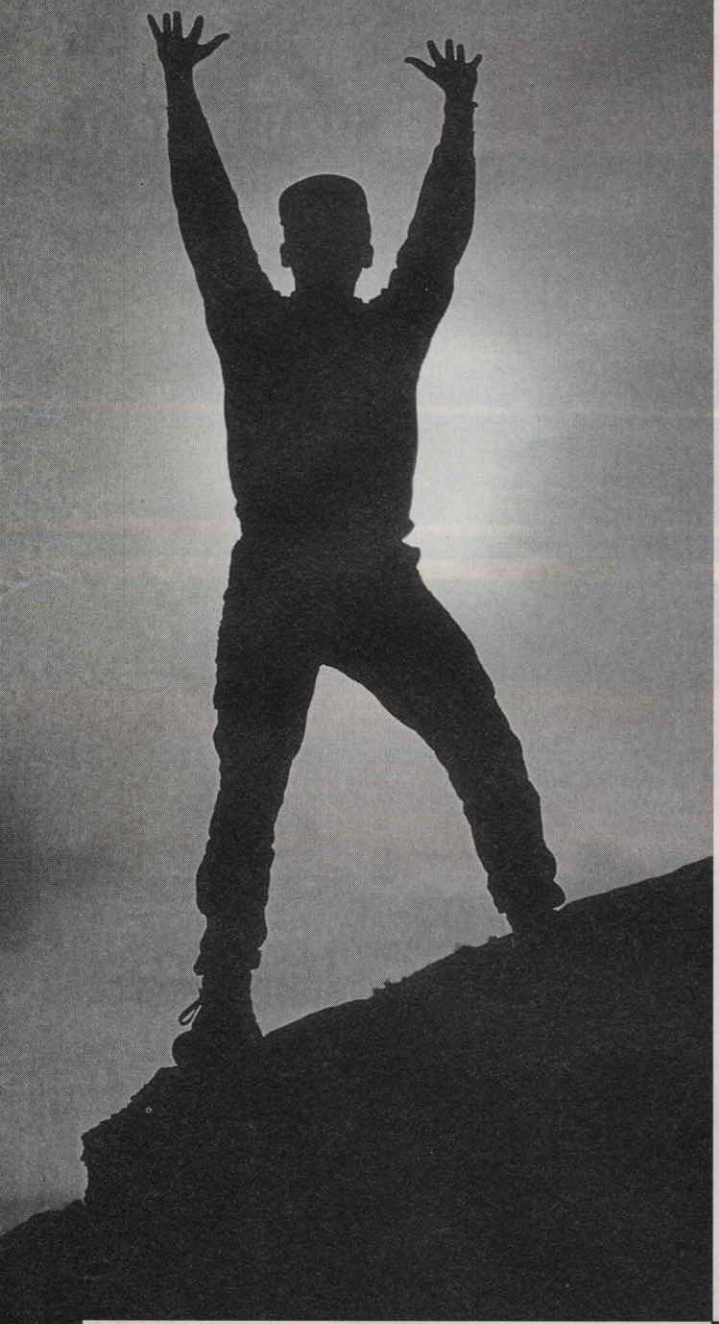
USA : (800) 507-2278
(415) 284-7970
Europe : (33) 1 46 56 88 47



CAST WORKBENCH™

by CAST
Move to the Forefront

"Gain Your Freedom"



NEW

Introducing CodeSQL™

"All the Speed, All the Power, None of the Royalties"

You Get a High Performance ODBC SQL Engine!

CodeSQL uses the ODBC standard to provide you with Oracle-level SQL capabilities. Deploy a wide variety of applications using the provided single-user, multi-user, and client/server configurations.

Use CodeSQL from C, C++, Visual Basic, Delphi, FoxPro, PowerBuilder, or any other front-end with ODBC support.

True Industrial Strength

CodeSQL conforms to the ANSI SQL Level 2, Oracle and ODBC Extended SQL grammar standards. For example, you can create/drop tables, views, indexes, and constraints. Inner and outer joins, nested queries, and other advanced SQL operations are all supported. You can even set isolation levels.

xBASE Compatible

CodeSQL uses xBASE files as its native file format. Built using CodeBase technology, CodeSQL is multi-user compatible with the data, index, and memo files of dBASE and FoxPro.

Nothing Moves Data Faster

CodeSQL is throughly optimized for performance. With our extensive query optimization, the more complex the query the more the relative performance increase. Gain even more speed through our automatic buffering which reduces network load, and with our statement caching which gives you static SQL performance with dynamic execution.

CodeSQL provides these additional benefits:

- Security
- Transaction Processing
- Full SQL Syntax Documentation
- Logging and Audit Trails
- 32-Bit Windows Support

FREE

With CodeSQL you get Royalty-Free Distribution!



Industrial Strength SQL

SEQUITER SOFTWARE INC.
P.O. Box 575 Newmarket NH 03857-0575
Fax: 403-436-2999
UK Tel: +44-181-317-4321
Internet: sequiter@supernet.ab.ca

Call Now To Order: 403-437-2410

Risk-Free, Unconditional 30-Day Money-Back Guarantee on all Sequiter product orders.

©1995 Sequiter Software, Inc. All rights reserved. CodeBase and Sequiter are trademarks of Sequiter Software Inc. All other product names are trademarks of their respective companies.

CIRCLE READER SERVICE NUMBER 9

Named Data Caches

Table Partitioning

Advanced SMP Configuration

Private Log Cache

Multiple Network Engines

Commit Sharing

Announcing

DBArtisan

OPTIMIZED FOR XI SYBASE SYSTEM 11



DBArtisan® 2.5



The leading administration tool for Sybase and Microsoft SQL Server just got a lot better. With DBArtisan 2.5, you can realize the dramatic performance gains of System 11 quickly. It simplifies the job of configuring data caches, partitioning tables and optimizing your SQL Server. DBArtisan 2.5 also incorporates strong **support for the major new features of Microsoft SQL Server 6.0.**

CALL TODAY FOR
A FREE 30 DAY EVALUATION
(415) 834-3131

Embarcadero Technologies
Tools for Client/Server Professionals™

MICROSOFT WINDOWS COMPATIBLE

SYBASE OPEN SOLUTIONS PARTNER

DBArtisan is a registered trademark of Embarcadero Technologies, Inc.
Other trademarks are the properties of their respective companies.

CIRCLE READER SERVICE NUMBER 251

*Open*Access™

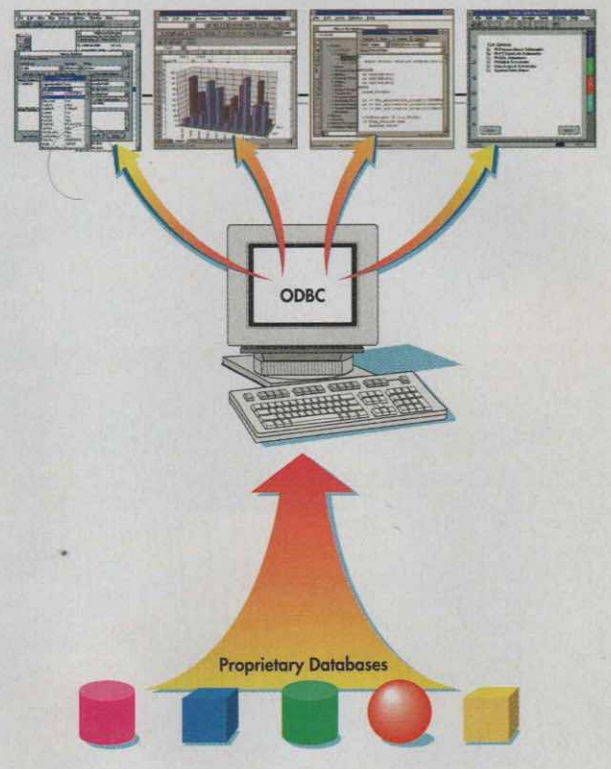
ODBC Driver Development Kit (DDK)

1. If you have valuable information locked in proprietary databases, *or*
2. If you need to access proprietary data,

*Open*Access ODBC DDK is for you!!

*Open*Access is *the* solution to distributed proprietary data access.

Have ODBC compliant access to your data source within a month!



The *Open*Access Solution Includes:

1. A robust ODBC driver tested with popular 4GL tools
2. Client/server connectivity to 11 different popular hardware platforms
3. An embedded SQL engine for parsing SQL request to a low level API, and
4. An easy to use toolkit for interfacing your specific data source

Applications of *Open*Access:

1. Providing ODBC access to non-SQL compliant systems
2. Accessing, retrieving, and storing data from proprietary systems in RDBMS data warehouse solutions
3. Universal GUI implementations for getting data from many specialized systems as well as relational sources using a single ODBC driver

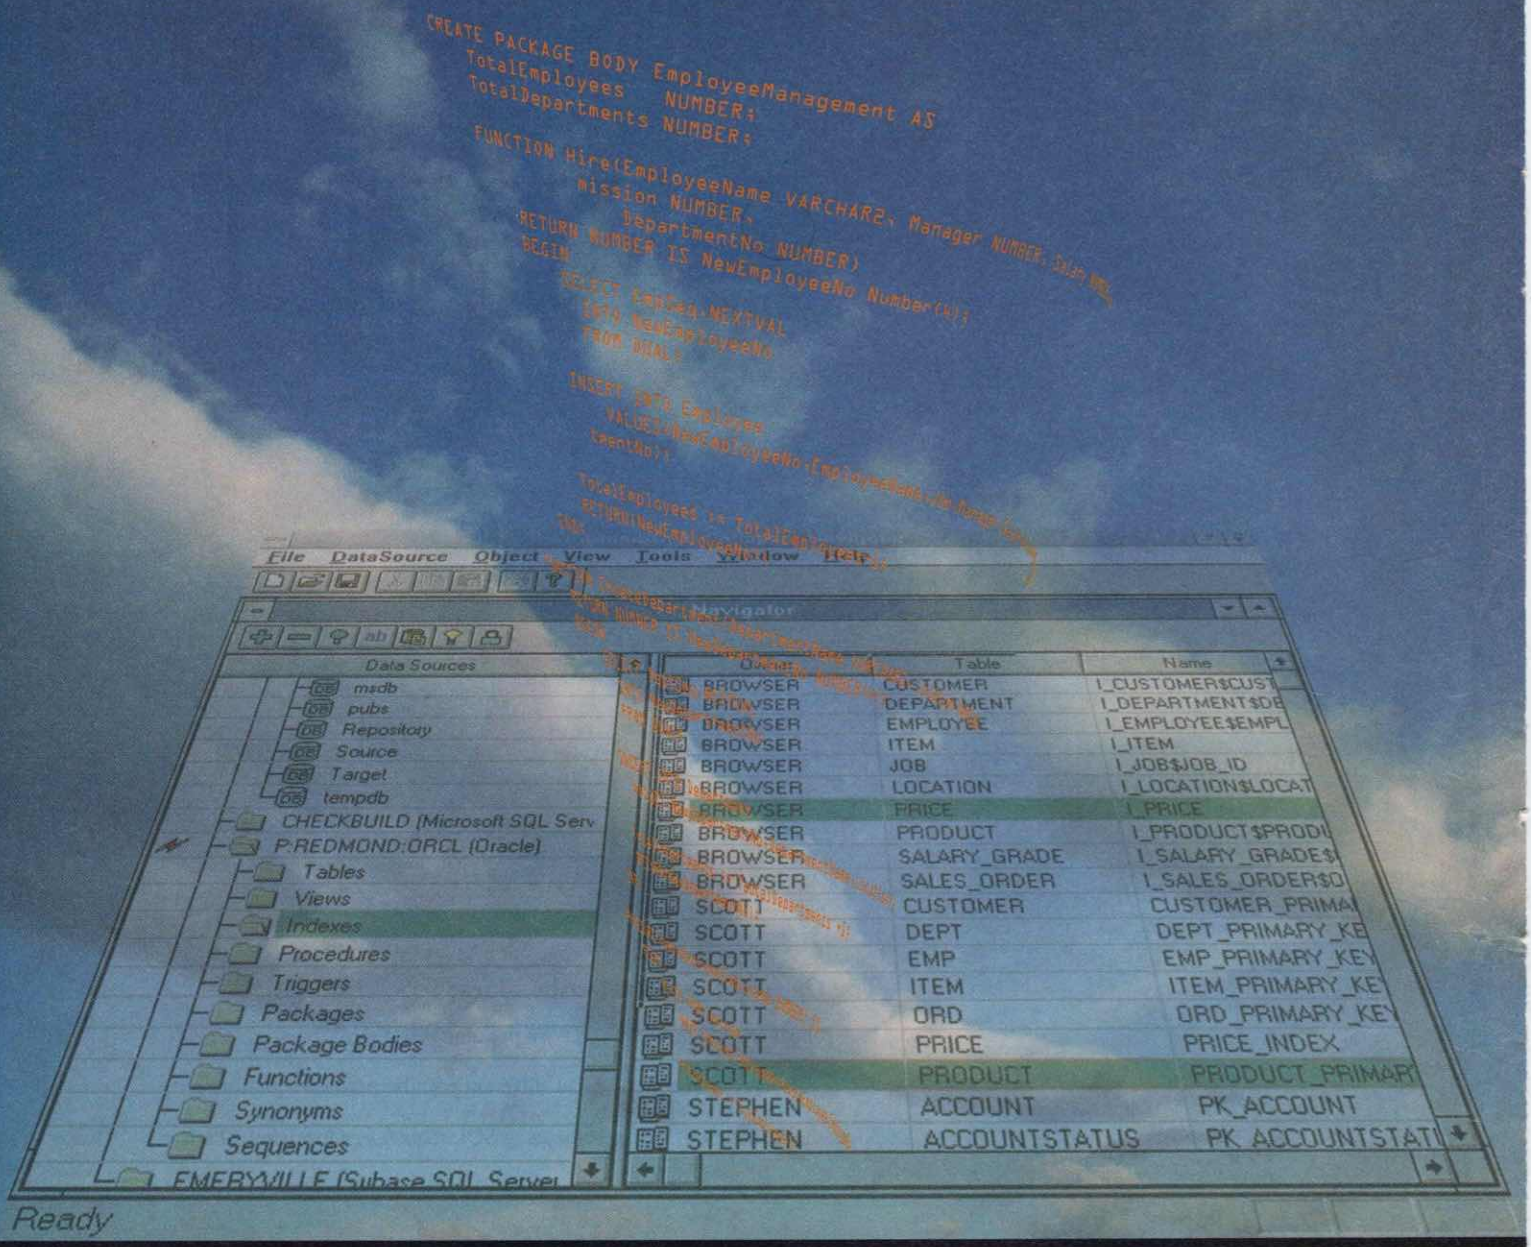
For Further Information, Contact:
Automation Technology, Inc.
1293 North First Street
San Jose, CA 95112
Tel: (408) 453-1099
FAX: (408) 453-8150
OpenAccess.info@atinet.com

ati automation technology, inc.

© Copyright 1995 Automation Technology, Inc. (ATI). All Rights Reserved.
*Open*Access is a trademark of Automation Technology, Inc.

CIRCLE READER SERVICE NUMBER 119

ORACLE 7 PROGRAMMING WITHOUT LIMITS



TeamSQL/Developer®

Your productivity will soar to new heights with TeamSQL/Developer, the first comprehensive development environment for Oracle SQL programmers. TeamSQL/Developer delivers a powerful scripting workbench, database navigation, reverse engineering and complete graphical object management in a Windows 3.1, Windows 95 or Windows NT programming environment. Get your project off the ground!

CALL TODAY FOR
A FREE 30 DAY EVALUATION

(415) 834-3131

$399
Competitive Upgrade
Limited Time Offer

Embarcadero Technologies
Tools for Client/Server Professionals™

Team SQL/Developer is a registered trademark of Embarcadero Technologies, Inc.
Other trademarks are the properties of their respective companies.

CIRCLE READER SERVICE NUMBER 117

SQL-PROGRAMMER®
for Windows™

ORACLE® version available!!! SQL-Programmer for Oracle Call today for details

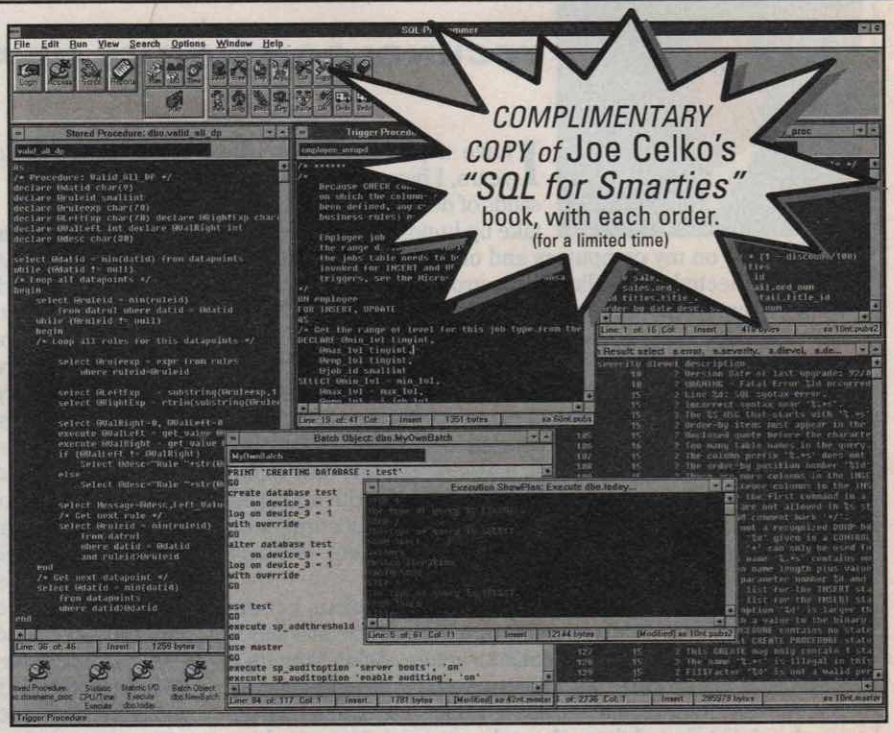
Sylvain FAUST INC.

SQL-Programmer is the most powerful, feature rich, Windows based, SQL Server back-end development environment aimed directly at increasing the programmer's productivity. It provides the Microsoft® SQL Server, SYBASE® SQL Server and Micro Decisionware Inc. Database Gateway developers with an intuitive environment for the development and maintenance of all programmable objects.

The Only Complete Development Environment which empowers the SQL Server programmer!

Some New Features ...

- Microsoft **SQL Server 6.0** Support
- **New Print Preview Feature** for reports, all printable objects and result sets.
- **Enhanced ISQL Compatible Script Generation** dialogue. SQL-Programmer will produce the DDL commands to recreate your entire database from devices right throught to the new SQL Server 6.0 objects (Operators, Alerts, Tasks, and Publications)
- **11 New Reports** have been added to provide complete documentation on your entire SQL Server 6.0 and below environment. SQL Programmer now boasts **a total of 27 different comprehensive reports.**
- **New Configuration Interface** allowing the SQL-Programmer user to customize the development environment.
- Automatic management of Administrative scripts via the ***NEW SQL-PROGRAMMER "BATCH" OBJECT***, managed by SQL-Programmer and stored within your desired SQL Server database (as a Central Repository). Fully integrated with the existing SQL-Programmer Check-In/Check-Out Team Programming Control.
- **New Editor workspace up to 16 Megs** per window, virtual size editing, including result set windows and FreeSQL window! Up to 255 levels of Undo and 255 levels of Redo per window
- **Editor faster than ever**



Features ...

- SQL-Programmer permits developers to **create and edit the programmable objects** so endearing to SYBASE and Microsoft SQL Server database engines.
- SQL-Programmer features an Industry unique **"VIRTUAL EDITOR"** which provides the developer with a **"test while you edit"** development environment for working on programmable objects like stored procedures and views.... It also provides **continuous foreground processing.**
- SQL-Programmer supports multi SQL Server connections.
- SQL-Programmer **extracts ISQL compatible scripts** from any of the connected databases and servers. Permitting the reproduction of the source commands used to define the entire database, or selected portions of it. Script generation can optionally produce INSERT statements for automatic repopulation of the table's DATA.
- SQL-Programmer supports team programming via the check-in check-out feature for programmable objects.
- SQL-Programmer recognizes the need for developers to be uninhibited and provides point and click access to a **"FreeSQL!"** window. In fact, multiple FreeSQL! windows can be opened at once on different servers/gateways and databases.
- SQL-Programmer automates the production of **execution plans, CPU usage and I/O statistics** directly from the SQL Server engine.
- SQL-Programmer features a **visual on-line search** within all programmable objects.
- SQL-Programmer, although extremely intuitive, includes both **on-line help, and printed documentation** plus automated direct access to HELPSQL.
- **And Much More...!!!**

To order, demos or for more information, call Ext. 152:
PH: (819) 778-5045 (fax on demand available)
FAX (819) 778-7943
Compuserve: 71162,1050 1-800-567-9127
Internet: info@sfi-software.com



SFI's family of programming interfaces for DB-Library™ & CT-Library™
Available as a VBX, OLE2 Automation Object or OCX (OLE Custom Control, VB4)

Introducing... SQL-Sombrero/OCX for OCI (Oracle Call Interface) Shipping Now!!!



Coming Soon!!!

Coming soon from SFI: "SQL- Agenda", a graphical Remote Batch Object Scheduler. This SQL-Programmer sister product will offer Microsoft SQL Server and Sybase SQL Server DBAs and developers a sophisticated means to schedule the execution of Batch Objects managed within their SQL Server environments.

Some Key Features are:

- Remote scheduling of SQL-Programmer Batch Object Execution
- Schedule the execution of "Bulk Copies" to and from your SQL Server
- Schedule the execution of Transact SQL scripts
- Controls scheduling remotely through a workstation's graphical user interface
- Manages SQL related scheduling for the enterprise across multiple servers

CALL FOR DETAILS TODAY, EXT. 152

© 1992-1995 Sylvain Faust Inc. Other product names herein have been used for identification purposes only, and may be trademarks and /or registered trademarks of their respective companies.

CIRCLE READER SERVICE NUMBER 73

BY TOM SPITZER

Needles in Document Haystacks



TEXT RETRIEVAL AND MANAGEMENT TECHNOLOGIES COME OF AGE.

I have a problem. Over the years, I have collected thousands of documents, all of which take up hundreds of megabytes of disk space on my computers and on the LAN to which I'm usually connected. My collection is growing dramatically as I surf the Internet and retrieve information on an ever expanding universe of topics, products, and technologies. I am able to find anything only because I trained myself to function as a human topic catalog early in life. When I was a student, I fancied myself a journalist and collected all manner of scraps of information. Before starting to collect information on disk, I already had sizable collections on paper, which was searchable only manually. Alas, as I enter middle age, I find that the breadth and depth of the information I and my colleagues are collecting is exceeding the capacity of my internal catalog system. We need tools.

There are several interesting facets to this problem. For one thing, it looks like a database problem, although the data fails to reveal itself in a particularly regular structure. Aside from its irregularity, a big problem with unstructured information is that it has very low information density. The facts I may need are dispersed and distributed throughout the material. Also, my documents are in a variety of formats, reflecting the document-creation tools that my company has used during the last decade. The distribution of documents across multiple file servers and Internet servers presents another challenge. Finally, I dream of a way to create a local index of interesting Internet sites that would enable me to return later when I need to see their contents. This would minimize the amount of document retrieval and local storage that I need to support my forays into cyberspace.

I realize that my problem is small compared to organizations, which retain much more information than I. My collection of information represents only a fraction of what my company has acquired in its 15-year life span. By most accounts, we are a relatively small company. Around the world, while billions of records exist in structured, field-oriented databases, the vast majority of corporate information is still created and stored in documents. The Gartner Group estimates that more than 90 percent of corporate intellectual property exists as text in various document formats. This corporate information base is a critical resource for knowledge workers, but cannot be accessed by most DBMSs.

Tom Spitzer is vice president of Advanced Products at SBT Accounting Systems in San Rafael, California. He is currently working on strategic development plans for the next generation of SBT management applications. Tom designed and managed the development of many of the products SBT has released since he joined the company in 1987. You can reach him via the Internet at tspi@sbt.com.

Full-Text Retrieval Systems

Where there are problems, there are usually candidate solutions, and I set out to find them. I discovered a highly competitive marketplace of full-text retrieval products. Full-text retrieval systems are designed to work well with large volumes of documents containing unstructured text and associated diagrams and images. An effective text-retrieval system provides the ability to store, manage, and subsequently retrieve textual information. The stored information can consist of complete documents, or it can be abstracts/bibliographies, email, notes, keywords, names, titles, and so on. Increasingly, these systems support multimedia information such as graphics, images, sound, and video as well. Effective systems allow users to access specific information without having preexisting knowledge of the content of the information base. More advanced text-retrieval systems allow nonexpert users to undertake searches in a natural, instinctive way.

A full-text retrieval system should allow you to enter data in an unstructured (free) format, as is typical in books, memos, and other text documents. Users must also be able to search and retrieve information based on the words and concepts in the documents, in conjunction with information recorded about the documents in structured fields. For example, a magazine article would also have an author name and a publication date. To fulfill these requirements, the core of these systems consists of very sophisticated indexing engines. It's interesting and a little amusing to find that many products claim to incorporate the most sophisticated indexing system on the planet.

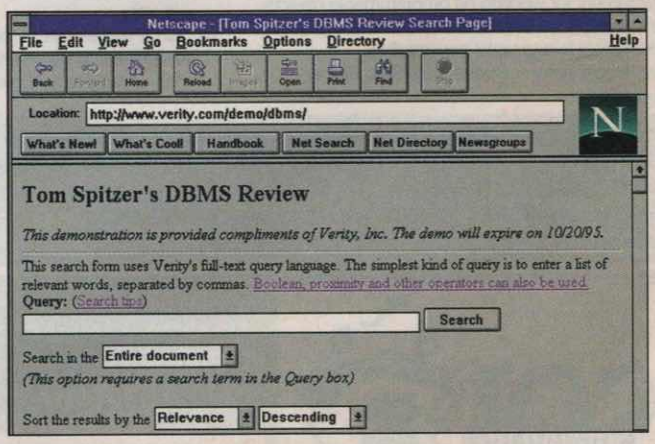


FIGURE 1 **An example of a Verity Search screen. This is a Web page that Verity built for me to search my *DBMS* articles.**

To determine how well a document fulfills a query, search engines measure the number of occurrences of search terms, the relative density of search terms in a document versus their density in all the documents in a search, the proximity of multiple search terms, and the location of search terms within a document. All of the products seem to have their own weighting algorithms. Some incorporate a user-modifiable thesaurus, which enables a search to find all documents that include the specified words and any words with the same meaning. Some products incorporate sophisticated stemming techniques for finding all occurrences of words with a common root. Xerox produces specialized lexical software for stemming, which other software vendors license. Probably the hottest technology in the field is concept or topic engines that allow users to specify concepts or topics, which the engines convert into word patterns. Most concept engines require user-defined concepts, which you can think of as saved queries. Newer technologies are being developed that allow concept searching without requiring you to first define concepts.

The products I discovered fell into several categories. I will describe two varieties this month, and leave some interesting emerging technologies for future consideration. The first category of products consists largely of programming APIs and indexing engines that developers can use to integrate document search and retrieval capabilities into corporate and commercial applications. The second category consists of full software products that provide the indexing and search engines as well as an environment in which the user works. Within each category, there are several variations. All of the vendors recognize the emergence of the Internet as a worldwide network of information resources and are developing products for Internet applications.

Verity

I researched this column almost entirely over the Internet. One of the first companies that I learned about was Verity Inc. Verity develops both tools and applications that enable intelligent search, filtering, and dissemination of textual information residing on enterprise networks, online services, the Internet, CD-ROM and other electronic media. Verity's Topic family of products includes Topic Enterprise Server, Topic Internet Server, Topic Client, Topic Agent Server, and Topic Developers Kit. Together they provide facilities for search, retrieval, and categorization of archived textual information, as well as real-time monitoring and filtering of incoming information such as news, alerts, or documents.

The Enterprise server indexes documents — regardless of their location or format — into a searchable Topic index, or "collection." The Topic Client is an attractive document navigation and browsing application whose key differentiating feature is agent technology. Topic Client provides "intelligent agents" that act on the user's behalf to search and retrieve specific types of information. Topic agents can sift through historic data (searcher agents), or you can deploy them as "watcher agents" looking for relevant information as it arrives. For example, Topic agents can watch data from a financial newswire (when coupled with the Topic News Server) to find information related to a user's stock portfolio. Investors can get timely results that help determine when to make portfolio decisions. The Topic Internet server adds the Topic engine's sophisticated full text indexing, search, and retrieval services to an existing Web site. These tools allow users to conduct personalized searches across Topic-indexed information stored within multiple sources and formats.

With the Topic engine, documents are not modified or altered for indexing. Documents are stored in their original locations and their native formats — always ready for retrieval, deletion, or modification upon request. The Topic engine supports several data formats, including ASCII, Adobe Acrobat's PDF, SGML, HTML, and more than 50 word processing, spreadsheet, and desktop publishing formats. The Developers Kit offers a suite of optional database gateways that provide access to documents stored in Oracle, Sybase, or ODBC-compliant databases.

I spoke to Michael Williams, Verity's product manager for the Internet Server product. He demonstrated the company's indexing and search engine by indexing my online collection of articles, and creating a search form for me to use to search them. The search form was an HTML page on Verity's server at http://www.verity.com. (See Figure 1, page 84.) The form was set up so that when I clicked the Search button, it would run a CGI script to call a program that searched the index of the articles at http://www.database.org. From the results, it dynamically generated an HTML page listing the matching articles, with hyperlinks pointing to them, so that when I selected one from the list, it would retrieve and display the full text of the article.

■ Topic Client provides "intelligent agents" that act on the user's behalf to search and retrieve specific types of information.

Verity is only seven years old, and was launched in large part to respond to the information management demands of both governmental and private intelligence organizations. That makes a great deal of sense. Such organizations have been collecting documents for hundreds of years. Recently the company has had substantial success licensing the Topic technology to prominent providers of document management and Internet products as well as to online services, including Adobe Systems Inc., IBM/Lotus (which incorporates it into Notes), Netscape Communications Corp., Quarterdeck Corp., and MCI's Delphi Internet.

Fulcrum

In a similar vein, Fulcrum Technologies provides tools for building custom applications for managing text-oriented organizational information. Fulcrum SearchServer features a scalable, client/server distributed processing architecture. Developers access SearchServer through interfaces such as SQL and ODBC, which the company hopes will enhance the product's long-term compatibility and interoperability with other system components. Independent and corporate developers can build custom applications to deliver SearchServer's advanced text searching and retrieval capabilities to their users. Such applications use a server API to talk to SearchServer with Fulcrum SearchSQL (SQL with text-retrieval extensions), SearchSQL statement processing, and results processing. Like Microsoft's ODBC, Fulcrum's API adopts the SQL Access Group and X/Open Call Level Interface (CLI) standards. For Windows, the SearchServer API is packaged as an ODBC driver, supporting the ODBC standard method of accessing data across multiple, heterogeneous data sources.

SearchServer organizes and references textual information through table structures. Each row in a SearchServer table corresponds to a document (or text object). You can optionally define columns in the table schema to record additional information not stored in the document itself. Unlike traditional database systems, you do not have to store the document text physically in the SearchServer table; SearchServer's text reader architecture allows organizations to keep documents in their existing locations in their original formats, such as Microsoft Word or WordPerfect files. Text objects stored in a relational database may be searched and retrieved transparently through SearchServer's database text reader architecture.

SearchServer builds a comprehensive index of terms to provide high-performance searching. During the search process, SearchServer accesses only the index files; the actual documents

are accessed only for display and indexing purposes. SearchServer supports incremental updating of its indexes. It gives you the option of updating the index in batch mode or whenever the table is modified.

The Fulcrum SearchBuilder products are toolkits for Visual Basic and PowerBuilder. These kits include custom visual controls and APIs for connecting to SearchServer and application templates. These toolkits enable developers to rapidly prototype, develop and deploy SearchServer-based custom applications using the same development tools they use for their more structured database applications. You can also connect to SearchServer via any ODBC-enabled application to build and maintain text-centric applications using off-the-shelf third-party tools.

PC-Grown Products

While the Verity and Fulcrum products are coming to the PC platform from Unix client/server backgrounds, several products come to text management from a PC background. Included in this group are some fairly venerable products, such as askSam and ZyIndex, both of which have had recent upgrades to enhance their usefulness in a Windows-based, Internet-oriented world. Also in this category is a product called Concordance, from Dataflight Software.

A "concordance" is an alphabetical index of all the words in a text or group of texts, showing every contextual occurrence of a word. With the Concordance software product, you create databases of documentary information and load them with documents, fragments of documents, or keywords with pointers to documents. Concordance tables are much like the tables you would create in a typical database product, with the added wrinkle that they provide a paragraph data type into which the database builder loads the textual data. After the database is loaded with documents, the index function builds a concordance of all the words stored in paragraph fields. When searching a Concordance database, you can create queries that combine field elements (for example, documents created after January 1, 1995) with text. Concordance has 20 search operators that are grouped into four functional areas: context, proximity, Boolean, and relational. The proximity operators *adj* and *near* allow the searcher to find words that occur up to 99 words apart. The Boolean operators include the operator *xor*, which is an exclusive-*or* operator that locates documents that include one but not both of the search terms.

ZyIndex, introduced in 1983, was the first PC-based full-text retrieval system. ZyIndex locates documents based on content, and provides an extensive toolkit for search customization. These include Word, Phrase, Boolean operators, Proximity, Quorum, Numeric or Date Range, Separators, and Wildcards. ZyIndex's progressive search feature lets you start a new search based on the results of the previous one. (See Figure 2.) Like Concordance, ZyIndex supports definition of data fields and compound searches based on a combination of fields and full-text information.

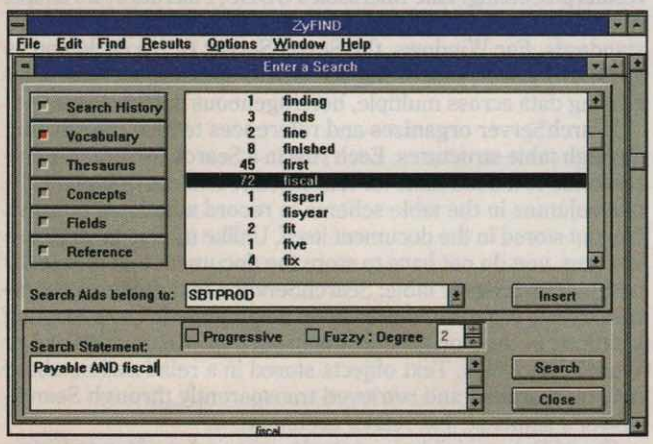


FIGURE 2 **An example of a ZyIndex Search screen. ZyIndex lets me search through the many documents I've collected about accounting topics.**

ZyIndex turns every word in your documents into keys that enable content-based retrieval. ZyIndex gives you search results within seconds, which can be ranked for relevancy based on user specifications. It leaves document files in place and searches through them in their native formats, with a wide range of supported formats, including WordPerfect, Microsoft Word, Lotus 1-2-3, Excel, and even PKZIP. After using ZyIndex to find a file, the user can open the file with the application program used to create it. In addition to making it convenient to make immediate modifications to source documents without ending a search session, ZyIndex provides an embedded hyperlinking feature for adding remarks, verbal notes, and video files to a source document without editing the actual text of that document. ZyIndex will replay any such annotations during subsequent viewing sessions.

ZyLab, the publisher of ZyIndex, is trying to catch the Internet wave with its ZyIndex for Internet product. ZyIndex for Internet is a Web server that runs on Microsoft Windows NT. When connected to the Internet, it lets users access full text indexes created with ZyImage or ZyIndex, using standard Web browsers. ZyIndex for Internet automatically generates HTML based on the user's query and related document content. Once the home page and a few simple templates are set up, Internet providers are spared from writing HTML. ZyIndex for Internet significantly speeds-up Web searching. Transactions are performed on a single index instead of on a series of sequential files. Selected information is converted to HTML in real time and passed to the client Web browser in the form of a list of relevant documents with pointers to the source files. This reduces bandwidth requirements for the client as well as download time.

askSam

Like ZyIndex, askSam has been around for a few years and has just released version 3.0 of its flagship product, a Windows version with expanded support for storing information collected from online services. An askSam user would create subject databases that combine both fixed-length fields and unstructured text components, and populate them by importing the documents created and collected from a variety of sources. The user could then search the database for combinations of words, words in proximity to one another, dates, or numbers. askSam also provides templates for importing email and forum messages from CompuServe filing cabinets into textual databases, an import filter for Lexis/Nexis stories, and the ability to use the product as both an HTML reader and authoring tool.

askSam has the most extensive feature set of the desktop products, offering many character and page formatting options, authoring tools, entry form and report design tools, mail merge, database statistics, and graphics file support. These features are all packaged in an attractive and fairly intuitive Windows product. I have used it for a while to manage my electronic correspondence database. Building a list of the 50 or so messages (out of more than 3000) that contain the string "DBMS" in either their subject or their body is virtually instantaneous. The list provides hyperlinks back to the source message so that I can immediately see the contents of any one of the matching messages. My only frustration with askSam now is that it does not interact more fully with a Web browser. When I click on a URL in askSam it displays the URL details in a dialog, which I can copy and paste into my browser (I'm using both Netscape and Microsoft Internet Explorer). It ought to save me the step by handing the URL off to the browser, or supplying an embedded browser of its own. On the receiving end, I have to save the site I'm browsing and then switch to askSam to import it.

I'd really like to see this work as a one-step process. For instance, askSam could provide a utility that hooks onto the browser and appends the Web site I'm browsing into an askSam database.

Future Tools

The increasing interest in text management has led to the entry of several high-tech software startups; I'll mention two briefly. The first is Architext Software. This is the brainchild of six recent Stanford computer science graduates who became frustrated with the current crop of text information management tools. These guys believed that: 1) most people searching for information don't know exactly what they are looking for, and 2) those who do don't know what keywords to use to find relevant documents. Their mission is to develop better retrieval tools and a suite of interactive browsing and hyperlinking tools that will help people navigate unfamiliar collections of data and quickly focus on interesting sources of information. Architext has attracted high-profile media attention as well as the interest of some of Silicon Valley's leading venture capitalists. I am currently using Architext's Internet search engine, called the Excite NetSearch (http://www.excite.com). Architext's concept search facility gives me the most relevant set of references of all the search engines I have tried. (See Figure 3.)

Another recent entry is LivePage from a company called Inforium, the Information Atrium Inc. Some well-known computer scientists at the University of Waterloo, Ontario, including a cofounder of Watcom and a former CEO of Waterloo Maple Software, came together to build a system of open, non-proprietary text and information management software products. They built their system on SGML and SQL standards; it stores SGML documents in popular relational DBMSs such as Oracle, Watcom, Sybase SQL Server, or Microsoft SQL Server. You can also use LivePage Tools to build custom database solutions with Visual Basic, PowerBuilder, and C. In the future, a standards-based approach that takes advantage of existing DBMS products and tools will make a lot of sense to organizations looking to implement text-retrieval applications.

An Evaluation Guide

I haven't mentioned all of the products that provide text-retrieval functions; however, my survey managed to uncover the key features and capabilities of these products and the directions in which their vendors are heading. As my opening comments implied, the ability to create indexes across diverse file formats and locations is important. Toolkits that let me build custom applications that integrate text-search functions into more traditional data management environments could be very useful. As the use of Web browsers for accessing information grows, the ability to search across Internet sites and format results into HTML documents will also be required. Finally, topic or concept search features may become far more prominent than the traditional formulation of search criteria as word combinations and operators. The Architext engine is just one example. I don't understand exactly how Architext maps my concept to an underlying word list, but I can see that this powerful search technology will be the foundation for the next generation of search engines. ■

- Architext Software, 2700 Garcia, Ste. 300, Mountain View, CA 94043; 415-934-3611 or fax 415-934-3610; Internet mail: info@atext.com; Excite NetSearch Web site: http://www.excite.com.
- askSam Systems, P.O. Box 1428, Perry, FL 32347; 800-800-1997 or fax 904-584-7481; Internet mail: emma@asksam.com; Web site: http://www.asksam.com.
- Dataflight Software Inc., 2337 Roscomare Rd., Ste. 11, Los Angeles, CA 90077; 310-471-3414 or fax 310-471-5294; Internet mail: info@dataflight.com; Web site: http://www.dataflight.com.
- Fulcrum Technologies Inc., 785 Carling Ave., Ottawa, Ontario, Canada K1S 5H4; 613-238-1761 or fax 613-238-7695; Web site: http://www.fulcrum.com.
- The Information Atrium Inc., 158 University Ave. West, Waterloo, Ontario, Canada N2L 3E9; 519-885-2181 or fax 519-746-7362; Internet mail: info@inforium.com; Web site: http://www.inforium.com.
- Verity Inc., 1550 Plymouth St., Mountain View, CA 94043; 415-960-7600 or fax 415-960-7698; Internet mail: info@verity.com; Web site: http://www.verity.com.
- ZyLab Corp., 19650 Club House Rd., Ste. 106, Gaithersburg, MD 20879; 800-544-6339, 301-590-2760, or fax 301-590-0903; Web site: http://www.zylab.com.

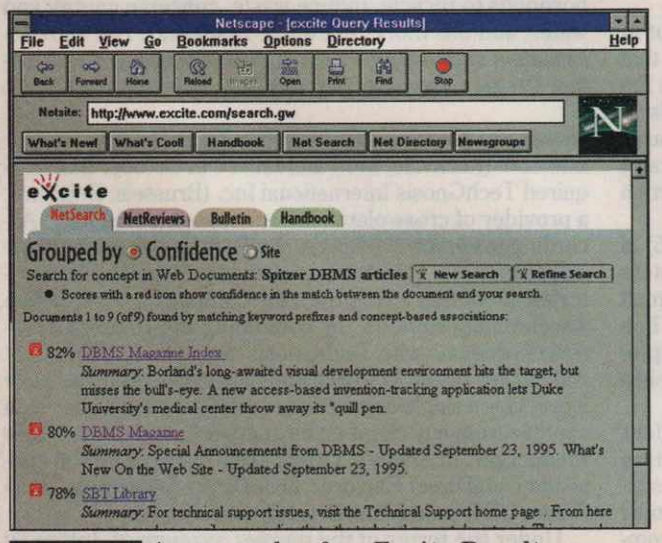


FIGURE 3 **An example of an Excite Results screen. Architext's Excite search engine let me enter a concept search for "Spitzer DBMS articles," and gave me better results than Lycos or Webcrawler.**

SQL EXPLORER *(continued from page 86)*

Actuals

```
item_no   actual_amount  check_no
==================================
10        300.00         '1111'
20        325.00         '2222'
20        100.00         '3333'
30        525.00         '1111'
```

Estimates

```
item_no   estimated_amount
==========================
10        300.00
10         50.00
20        325.00
20        110.00
40         25.00
```

Your problem is to get this output from a single query :

Results

```
item_no  description  tot_act  tot_est  check_no
=================================================
10       'item 10'   300.00   350.00   '1111'
20       'Item 20'   425.00   435.00   'Mixed'
30       'Item 30'   525.00   NULL     '1111'
40       'Item 40'   525.00   NULL     NULL
```

Item 50 from the Items table is not to be shown because there was no record for it in either the Actuals or the Estimates table. The column tot_act is the total of actual amounts for that item; the column tot_est is the total of estimated amounts for that item. (Puzzle answer on page 25.) ■

■ SYSTEM 11 DEBUTS
■ MAGIC DEVELOPERS COMPETITION
■ INTERSOLV BEEFS UP ITS MIDDLEWARE

Sybase Execs Preview System 11

On October 30, 1995, Sybase Inc. (Emeryville, Calif.) formally announced Sybase System 11, a suite of tools that includes the Enterprise Control server toolset, the Enterprise Connect middleware and gateway toolset, and the Powersoft toolset. To complicate the mix further, Sybase has renamed several of its tools to "strengthen their positions in the System 11 strategy." Enterprise Control is composed of Sybase SQL Server, Sybase IQ (formerly IQ Accelerator), Sybase MPP (formerly Sybase Navigation Server), and Sybase SQL Anywhere (formerly Watcom SQL). Enterprise Connect is composed of Replication Server, DB Gateways (formerly MDI Gateway), Open Server, OmniConnect, and Enterprise Messaging Services. The Powersoft family is composed of PowerBuilder Enterprise and Desktop, the recently acquired S-Designor modeling tool, InfoMaker, and the Watcom compilers.

In a lavish presentation in New York (broadcast via satellite to West Coast press and analysts), Sybase Chairman, President, and CEO Mark Hoffman said that System 11 should lay to rest "any lingering questions of performance, scalability, and quality." Using a triathlon metaphor, company officials hammered on this point, and had three of its biggest customers give testimonials on how System 11 performs well, scales well, and provides accurate results. The System 11 strategy provides a single architecture from which customers can choose a total solution or individual products.

Sybase Senior Vice President of the Enterprise Products Division Dennis McEvoy outlined SQL Server 11 (due out in December 1995), which features a Logical Memory Manager optimized for mixed-workload environments. This Logical Memory Manager provides named caches, variable block sizes and buffer algorithms for faster I/O, a cache-sensitive optimizer, and a cache analysis tool. SQL Server 11 will also feature dynamic load balancing, and it is ISO 9000-certified for quality. (Dennis McEvoy provided a detailed preview of Sybase System 11 in *DBMS*'s October issue, page 38.)

In the data warehousing space, Sybase IQ (due out in the first quarter of 1996) is an extension of SQL Server that will feature bitmap indexing technology (acquired from Expressway in early 1995). Sybase MPP (available now) is an optional add-on to SQL Server that is optimized for parallel computing. According to McEvoy, Sybase MPP scales linearly to 128 processors and is in use now for production databases with 450GBs of data and 300 users.

Sybase SQL Anywhere (due out in December 1995) is positioned for mass deployment, from mobile computers to the desktop to the LAN. It offers low cost with a small footprint, and Sybase Transact-SQL compatibility for better integration with SQL Server. Plus, SQL Anywhere provides two-way replication to/from remote and mobile sites via SQL Remote.

Powersoft CEO and Sybase Executive Vice President Mitchell Kertzman provided a rather vague view of future Sybase plans, saying only that Sybase will continue to support new technologies, including the Internet and electronic commerce, multitier applications, and network-centric computing. While Kertzman said that "System 11 provides a great foundation for the next generation of computing," he did not give any specifics or name future products. *Sybase Inc., 510-922-3500.*

1995 Developers Competition Winners

In early October, 50 teams of independent and vendor-sponsored developers traveled to Durham, N.C. for the 1995 Developers Competition, produced by Droege Computing Services Inc. and Advisor Publications. Each team had to produce a fully functional application in one day, based on a predefined specification and late-breaking modifications. The first five overall winners were as follows (special congratulations to Magic Software Enterprises [Irvine, Calif.] for sweeping the first five places):

- *First Place:* James Holcombe (Magic Software Enterprises, Norcross, Ga.) and Mike Rose (Health Communications Services, Norcross, Ga.), using Magic, CTree, Oracle, and Btrieve; developed in DOS and ported to Windows.
- *Second Place:* Bruce Lomasky (Magic Software, W. Hartford, Conn.), using Magic; developed in DOS and ported to Windows.
- *Third Place:* Keith Canniff and Sherman Levine (Management of Information Technologies, Norcross, Ga.), using Magic; developed in DOS.
- *Fourth Place:* Cindy Webb and Meeta Ajjampur (Structured Computer Systems, Avon, Conn.), using Magic; developed in DOS.
- *Fifth Place:* Heidi Schuppenhauer (Trillium Software, Snohomish, Wash.) and Denise Schillaci (The Resource Group, Snohomish, Wash.), using Magic; developed in Windows.

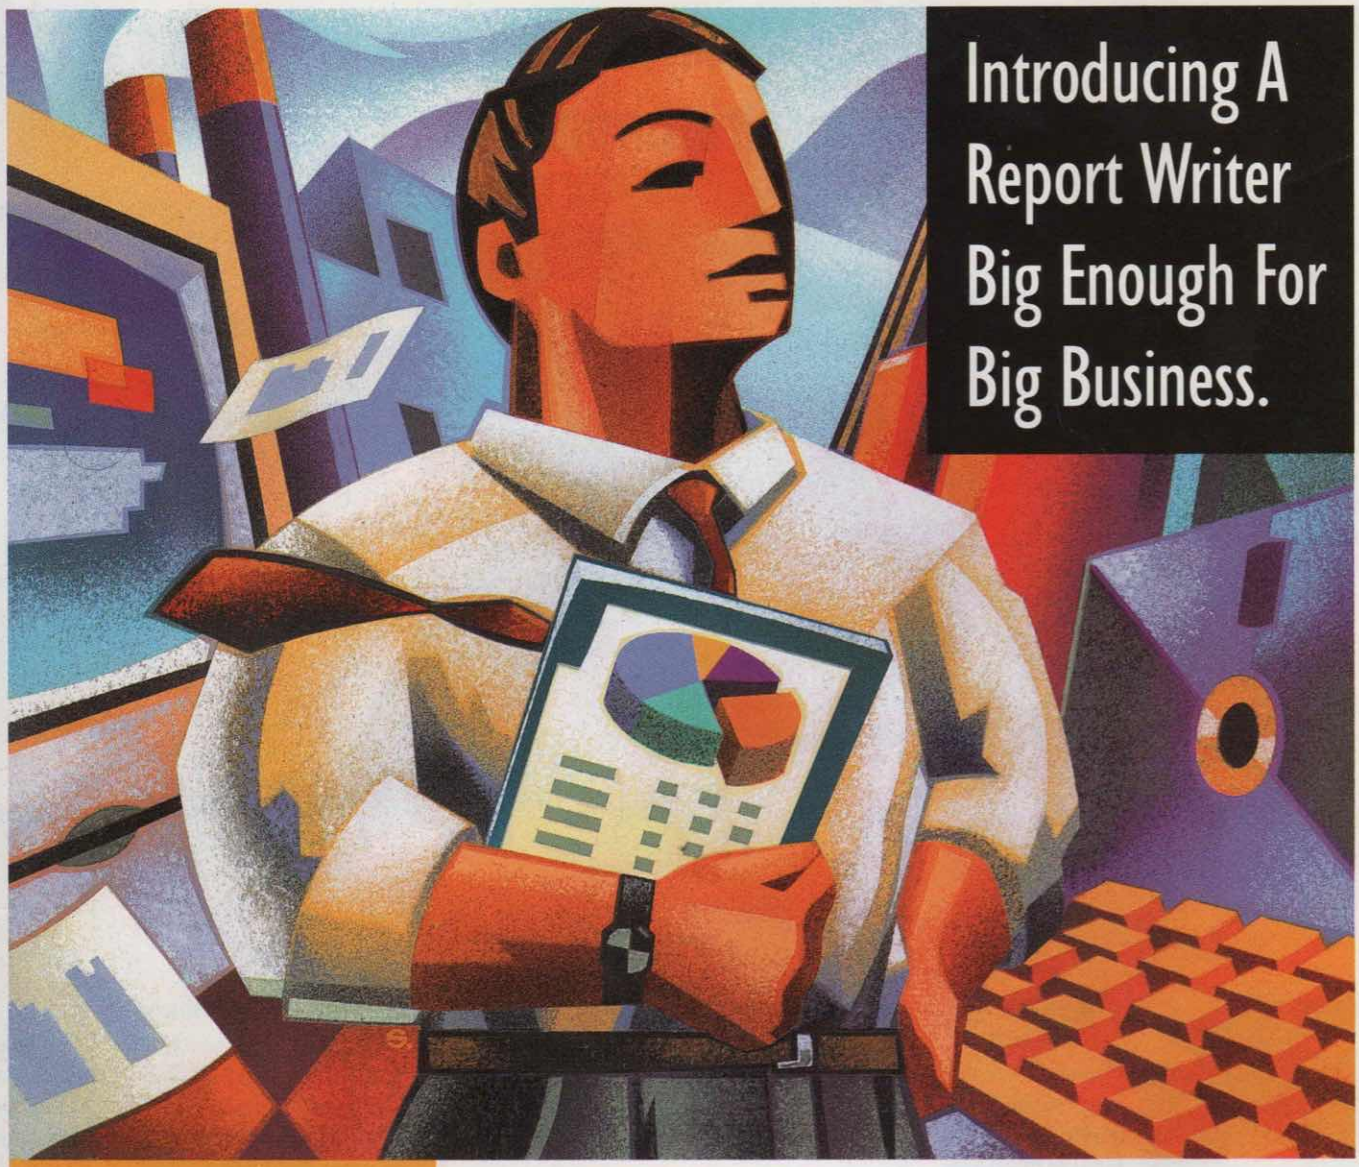
Bronze, silver, and gold medals were also awarded in special categories. Client/server gold medals went to developers using Oracle Developer/2000; GUI gold medals went to developers using Magic, CTree, Oracle, and Btrieve; OOP gold medals went to developers using Gupta SQLWindows; and RAD gold medals went to developers using Magic.

The applications were designed for SunShares, whose mission is to "help people use the earth's resources in a more sustainable, efficient, and healthful manner." SunShares provides opportunities for communities and neighborhoods to recycle, reduce waste, conserve energy and water, and use more renewable resources. *For more information on SunShares or the Developers Competition, contact Droege Computing Services Inc., 919-383-9749.*

Intersolv Comes Down with TechGnosis

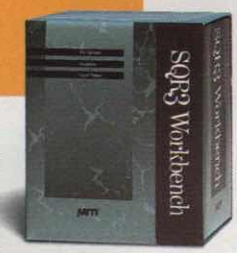
Intersolv (Rockville, Md.) announced in October that it acquired TechGnosis International Inc. (Brussels, Belgium), a provider of cross-platform data access technology. According to company officials, the TechGnosis product line will strengthen Intersolv's presence in Asia/Pacific territories, as well as strengthen the company's "virtual data warehousing solution" by combining the DataDirect series of products with TechGnosis' SequeLink. SequeLink will bolster Intersolv's back-end data access tools because it provides a single connection to relational and legacy data sources from any client on any network or platform. After the merger, Intersolv's Virtual Data Warehouse will consist of DataDirect Explorer, DataDirect SmartData, DataDirect ODBC, and now DataDirect SequeLink.

Under the terms of the merger agreement, Intersolv acquires TechGnosis for approximately 3.4 million shares of Intersolv common stock and $7.5 million in cash. The transaction is valued at close to $80 million. *Intersolv, 301-838-5000.* ■



Introducing A Report Writer Big Enough For Big Business.

SQR3 Workbench™ The Universal Production Report Writer.



The heavy-duty, industrial-strength reporting challenges of today's large company environments have created a world-wide, red-eyed phenomenon: long hours and code-crunching nights of sweat and frustration as MIS departments try to wrestle more productivity from their client/server technology.

Think about a more productive approach: SQR3 Workbench. More than a report writer – a production report writer. The only suite of solutions geared for complex, high volume, enterprise-oriented report writing.

Where other report writers are application specific, SQR3 Workbench operates with any database, across virtually any platform, at any level of complexity on either the client or the server. Where other third party products choke on too much data, SQR3 Workbench can handle infinite amounts of information through data paths of literally any complexity. And where others constrain the format and portability of the output, SQR3 Workbench processes reports or procedures without restriction and with dramatic time savings.

Based on the established and popular SQR,® with over a million users worldwide, SQR3 Workbench is an elegant procedural extension of full ANSI SQL, with 4GL constructs. It ships complete with Easy SQR3,™ a 4GL interface and quick report prototyping tool; SQR3 Developer's Kit for debugging; and SQR3 Execute™ for deploying runtime applications. SQR3 Viewer™ a Windows™-based broadcast report viewer, can be distributed across the enterprise, free.

So no matter how big the enterprise, no matter how big the report writing challenge, now you have an approach that means business. Get busy. For more information regarding SQR3 Workbench or to discuss your particular requirements, call one of our Technology Sales Representatives at **1-800-366-6484, Ext. 361**; or FAX at **310-424-9385**.

MITI®

© 1995 MITI. MITI and SQR are registered trademarks of MITI. All other brand names are trademarks or registered trademarks of their respective holders.

CIRCLE READER SERVICE NUMBER 68

New Paradox 7 for Windows 95 and Windows NT

Adapt to your environment effortlessly.

Store a wide range of data types from number, text, date, and time to graphics and multimedia.

Build all levels of business solutions from single-user PC-based applications to LAN-based or client/server applications.

Use OLE Automation and OCX controls to make Paradox work seamlessly with other Windows applications.



View and organize your data with Windows 95–hosted features such as new tab-style controls in the Visual Form Designer.

Create customized forms displaying data from multiple tables with visual design tools—no programming required.

Customize your working environment—see the toolbar how and where you want it.

"What *InfoWorld* has, for a long time, considered to be the hands-down best Windows database… is now more powerful than ever as a development environment."

—*InfoWorld*, October 16,1995



Evolution on demand

Paradox® 7 will help your transition to a whole new world. Find out for yourself why *InfoWorld* says, "This is the first Windows 95 application that is improved enough to be a compelling reason to upgrade from Windows 3.x to Windows 95 and Windows NT."

What's New

New! Enhanced OLE support lets you categorize, control, manage, store, and share virtually any type of data—including documents, spreadsheets, and graphics.

New! OCX support means that developers can now create applications that use OLE custom controls such as Internet Web browsers, communications packages, or grid and graphing controls that seamlessly integrate with Paradox.

New! More than 100 new ObjectPAL methods give you the most powerful database language available.

New! Object Explorer helps you accomplish more by accessing properties, methods, and events, no matter what your level of expertise.

Unparalleled compatibility

With enhanced Experts and a new Windows 95 look and feel, you'll see the benefits of the new 32-bit technology from day one. Paradox 7 is compatible with previous versions of Paradox for Windows, so any data or applications you created in Paradox for Windows can be used in Paradox 7. You can even share your data with other people who may still be using previous versions of Paradox for Windows on the same network server.

Best suite integration

With new act-alike menus and toolbars, you'll appreciate the seamless integration of Paradox with Novell PerfectOffice 7 and Microsoft Office 95.

Powerful development language

Paradox 7's Integrated Development Environment is unsurpassed. New developer tools such as the enhanced program editor make ObjectPAL™ programming a snap. The editor boosts productivity with features such as color syntax highlighting, block indent, keystroke record and playback, and other tools. Our new high-performance 32-bit Borland® Database Engine (BDE) technology is at the heart of all Borland products. That means when you use Paradox 7, you can access more diverse data faster and easier than with any other database.

Best network solution

If you use Paradox on a network, you'll appreciate its LAN support. With its record-locking capabilities, Paradox allows you to create high-performance network applications, making it your best networkable database choice.

"It is without a doubt a must-have for all serious Paradox users and developers."

—*InfoWorld*, October 16,1995

Try Paradox 7 risk-free

If for any reason you're not fully satisfied, simply return it to Borland within 90 days for a full refund.

Order Paradox 7 Today

FOR ONLY **$89.95***

See your reseller or call today.
1-800-336-6464, ext. 50389

To see the entire *InfoWorld* article, just access our web site.

Canada: 1-800-461-3327 • Internet: http://www.borland.com
CompuServe: GO BORLAND

Borland
Making Development Easier

*Limited-time offer after manufacturer's rebate. Copyright © 1995 Borland International, Inc. All rights reserved. All Borland product names are trademarks of Borland International, Inc.
Offer good in the U.S. and Canada only. All prices are in U.S. dollars. Dealer prices may vary. BI 8709

CIRCLE READER SERVICE NUMBER 135