INSIDE! Spectrum 2004 Review: Find Out About the Latest Technological Innovations the Industry Has to Offer

$7.00 U.S.

INTERNATIONAL

SPECTRUM®

THE BUSINESS COMPUTER MAGAZINE

MAY/JUNE 2004 • AN IDBMA, INC. PUBLICATION

History of MultiValue

Where Are They Now?

The Original Microdata Dealers: A Fascinating Story of Greed, Competition, Camaraderie, and Wealth

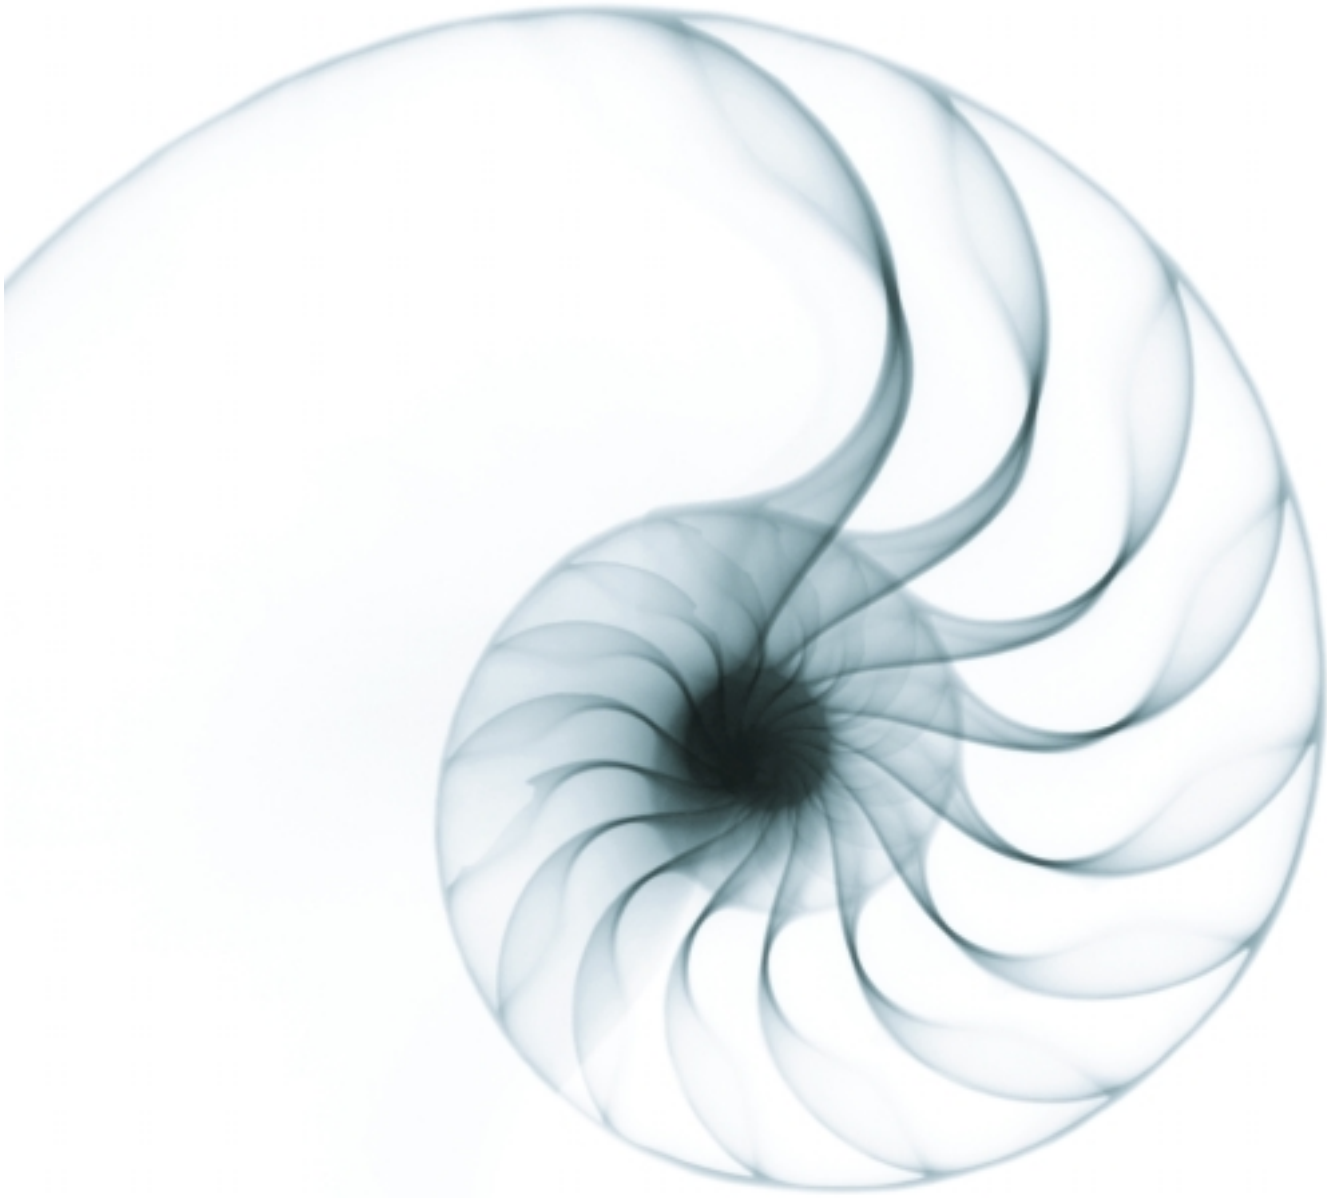

"THEY BELIEVE DESIGNBAIS WILL....SIGNIFICANTLY IMPROVE YOUR WEB DEVELOPMENT PRODUCTIVITY, SIGNIFICANTLY IMPROVE YOUR PROFITABILITY, IMMEDIATELY TURN YOUR MULTI-VALUE DEVELOPERS INTO WEB DEVELOPERS"

WILL IT REALLY DO ALL THAT?

YES

DesignBais

www.designbais.com

Cover Story

HISTORY OF MULTIVALUE: THE ORIGINAL PICK DEALER NETWORK — WHERE ARE THEY NOW?

16 Huge opportunity, greed, competitiveness and camaraderie: All of these are ingredients in the behind-the-scenes story of the rise and fall of the Microdata Dealer Network.

DEPARTMENTS

INTERNATIONAL SPECTRUM®
MAY/JUNE 2004

THE International MultiValue Marketplace

Up Close and Personal, PART 1

An International Affair

INTERNATIONAL SPECTRUM 2004, held in Las Vegas, Nev., was the gathering place that brought together contingents of MultiValue users, database companies, vendors, developers and VARs from around the world. With representatives from Australia, the United Kingdom, and other parts of the world, it was truly an international affair.

Underscoring the growth and popularity of the MultiValue market across the globe, "ambassadors" from Australia, Martin Meier, president of Meier Business Systems (MBS), and from the U.K., Alf Pilgrim, chief technology officer of Northgate Information Solutions, made remarks at the MultiValue breakfast co-sponsored by Raining Data and Microsoft.

Meier remarked that "PICK is not dead. I see us as very actively supporting MultiValue products, keeping them alive. Demographically, in Australia, there is a particularly high installed base of MultiValue systems across every vertical."

He shared that MBS has experienced 18 years of year-on-year growth. "Many of my peers in the MultiValue market are equally successful," he said. "MBS is 110 percent committed. The economy is good Down Under, and it's a great place to do business."

As a representative from the U.K., Pilgrim commented, "There is a thriving European and U.K. MultiValue market that is fairly starved for new MultiValue information and products."

Margins are typically higher in the U.K., he said, as well as its culture being more accepting of products that don't fit a particular mold. He encouraged companies to come to the next International Spectrum show scheduled in London on September 23-24, 2004. "The U.K. is a stepping stone to Europe [for businesses]; it's a good place to get a hold, then step into the rest of Europe."

The opportunity is there for MultiValue businesses to capitalize on failed attempts by companies who've tried to change their IT systems, but haven't had the success they hoped for, Pilgrim said. "Many companies have lost millions implementing huge IT projects," he said. "Other projects are successful but have extensively high support costs. Instead of freeing them, they've been hamstrung by that. They want answers, and they're not going to continue to pay these support costs. They want a long-term, sustainable solution."

Urging the MultiValue community to "hold the flag for MultiValue," Pilgrim commented: "I believe all of us in this room believe that MultiValue has something to offer. I'm sick and tired of the oft-repeated and hackneyed phrase 'industry's best kept secret.' I often heard at the U.K. show [held September 24 and 25, 2003] from people who felt the show had uplifted them and given them a positive message to pass on to their companies."

MultiValue Ushers in New Technologies With a Flourish

On the show floor, there were demonstrations that definitively proved that the MultiValue environment is primed and ready to interoperate with leading technologies like Microsoft .NET, Oracle, and XML. The products that were demonstrated were a true testament to the validity of the MultiValue model and the thousands of business applications that have been built around it. As Susie Siegesmund, director of U2 Business, IBM Data Management, stated at the show: "I'm sure everyone realizes that the world is becoming more cognizant of the MultiValue model of thinking. My message is, 'Don't give up the ship, it's the greatest technology; it's the greatest to develop in; and it's the most flexible model you can hope for in the world today."

Read on to see what MultiValue vendors are up to these days.

Continues on page 8

That's why we've dedicated our 2004/2005 conferences to showcasing the newest and best software tools available for you to take your MultiValue application to where you want it to go! **Come see what's new!**

For more information visit **www.intl-spectrum.com**

London, England
September 22 - 23, 2004
Novotel London West Hotel

Sydney, Australia
October 12 - 13, 2004
Star City Hotel & Casino

Melbourne, Australia
October 18, 2004
La Meridien at Rialto

San Diego, CA, USA
February 22 - 25, 2005
Hyatt Regency Islandia
Hotel & Marina

IDBMA, Inc.'s
INTERNATIONAL
Spectrum®

[FROM THE **INSIDE**]

The Looming War Over SPAM

On March 15th, *International Spectrum* magazine launched Spectrum *e-Xtra*, a new "e-newspaper" that is sent out via email on the 1st and the 15th of each month. If you're not receiving *e-Xtra* and you'd like to, just update your subscription information on **www.intl-spectrum.com**.

Since *e-Xtra* is just one of the new genre of emailed newsletters, the topic of email in general heads the editorial list. In the April 1st edition, for example, the new USA federal CAN-SPAM Act of 2003 was covered and how it was rushed into law to preempt punitive state laws before they could go into effect.

Like all other email, people either like *e-Xtra* or they consider it to be spam. One irate recipient vowed to put us on every blacklist in the world — probably, he was having an extremely bad day. Another moaned that the newsletter was 280K — what's he running on — a 386? One guy didn't want to receive the newsletter, but he wanted to get all of our other emails. Right! Let me just set up the mechanism to handle that right away!

But that's O.K. — we always remove someone who wants to opt out, or try to accommodate special requests when we can. But the one thing we can't do anything about is the person who would like to receive *e-Xtra*, but his local provider or the company firewall won't let it through. For that matter, many disappointed people can't even receive simple text emails from people they want to communicate with for the very same reason.

This, in my opinion, is the looming war over spam which will change the entire face of the Internet in the future. Who gets to decide if something is spam? What are the liabilities of blacklisting somebody? What do you mean I can't receive emails from my customer? What do you mean I can't send my customer his order status? How often have you had this conversation lately? "I don't understand — I sent it to you last week!"

I really feel for the lawmakers who are going to have to sort this out. How do they get rid of the mega-spammers and the porn people without trampling on legitimate business. Who should be the policeman — the Internet Service Providers? What do they do about the fact that the Internet is an international phenomenon? How long will it be before the class-action suits begin over innocent companies that are destroyed by blacklisting?

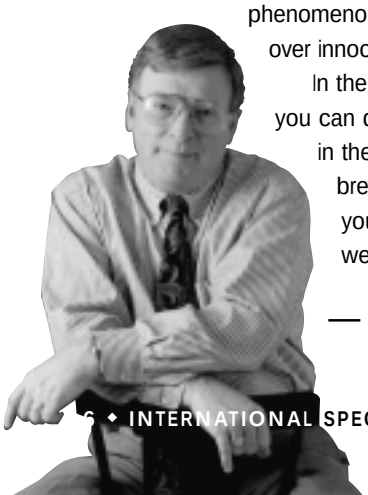

In the meantime, if you want to get all of your relevant email, you can do what I did, which is to disable the anti-spam software in the firewall and sift through the penis enlargements and breast enhancements to get that email from Aunt Susie. If your Internet Service Provider is filtering your email for you — well, that's a problem!

— GUS GIOBBI, CHAIRMAN, IDBMA, INC. —
gus@intl-spectrum.com

INTERNATIONAL
SPECTRUM®
MAY/JUNE 2004

MISSION STATEMENT *International Spectrum* magazine's editorial mission is to be the premier independent source of useful information for users, developers, and resellers of MultiValue database management systems, open systems business database solutions, and related hardware, software, and peripherals. Published bimonthly, *International Spectrum* provides comprehensive coverage of the products, companies, and trends that shape the MultiValue marketplace as well as the computer industry at large — helping its readers get the most out of their business computer systems.

IDBMA's *International Spectrum* is published six (6) times per year at the subscription price of $40.00 U.S. in the U.S.A.; $45.00 U.S. in Canada and Mexico; $50.00 U.S. for other countries. Single copy rates are $7.00 U.S. in the U.S.A. and Canada, and $9.00 U.S. in all other countries. *International Spectrum* is published by IDBMA, Inc., 311 4th Avenue Suite #513; San Diego, CA 92101; Tel: 619/515-9930; Fax: 619/515-9933 **E-Mail:** requests@intl-spectrum.com; **Website:** http://www.intl-spectrum.com.

PRINTED IN USA • ART AND DESIGN: CP Design; 858-642-6878; San Diego, CA

NEWS RELEASES/UNSOLICITED ARTICLES
International Spectrum is eager to print your submissions of up-to-the-minute news and feature stories complementary to the MultiValue marketplace. Black and white or color photographs are welcome. Although there is no guarantee a submitted article will be published, every article will be considered. *International Spectrum* retains all reprint rights.

PUT YOUR MONEY ON A SURE THING.

OPENINSIGHT
Database Development

JOI
Java for OpenInsight

Using Revelation Software's flagship products OpenInsight and Java for OpenInsight (JOI) are truly safe bets! OpenInsight is *the application developer's toolkit.* It has its own development environment and database, including a Form Designer, Editor, User Interface Workspace, Reporting Tools, built-in Web and XML functionality, plus much, much more.

If you need a robust, secure environment to develop applications on a Linux or Windows platform, then you need the power of Revelation Software.

Download a free 30-Day Evaluation Copy of OpenInsight or JOI today. It'll be like hitting the jackpot.

For more information about OpenInsight, our Solution Partner Program, or OpenInsight training classes at Revelation Headquarters or one of our Authorized Training Centers (ATCs), go to www.revelation.com/intl-spectrum/ishome or give us a call at 800-262-4747

2004 Revelation Software User's Conference in New Orleans

Join us at the Revelation Software User's Conference in New Orleans, Louisiana, June 24-26, 2004. Bring this ad for a free gift. Go to www.revelationconference.com for more details.

Revelation Offices – Global

Revelation Software, Inc.
Corporate Headquarters
99 Kinderkamack Road, 1st Floor
Westwood, NJ 07675
Toll Free: (800) 262-4747
Phone: (201) 594-1422
Fax: (201) 722-9815
Email: info@revelation.com
Web: www.revelation.com

Revelation Software Ltd.
Northumberland House
11 The Pavement
Popes Lane, Ealing
London, W5 4NG, UK
Phone: +44(0)208 832 7380
Fax: +44(0)208 832 7381
Email: info@revsoft.co.uk
Web: www.revsoft.co.uk

Revelation Software Australia Pty Ltd.
Unit 2
21 Old Pittwater Rd.
Brookvale, NSW 2100
Australia
Phone: +61-2-9939-6399
Fax: +61-2-9939-6366
Email: info@revelationsoftware.com.au
Web: www.revelationsoftware.com.au

REVELATION SOFTWARE

Candid Comments *at Spectr*

IBM: Committed to U2

Since IBM acquired the UniData and UniVerse database management systems (U2) in 2001, the question of whether IBM will continue to develop and support the U2 products just keeps popping up, fueled by rumors despite IBM's statements that it will continue its support.

At Spectrum 2004, IBM's Susie Siegesmund succinctly quelled the rumors once again: "IBM remains completely committed to U2." She further explained where IBM's strategy is headed. "The neat thing from our viewpoint is that IBM has moved into focusing on the software group level. The software group has five pillars—one is data management, which U2 is a part of. It's exciting for us because this last year we've reorganized the sales and marketing teams at the software group level. The idea is solutions sales, so we're really focusing on sales through partners. Of course, that plays very well for my partners because that's how we sell the U2 product — through our partners.

A conversation at the IBM booth

"That's also exciting for us as a team because suddenly we're not seen as just one of several databases — and in truth, not viewed as the primary database that IBM delivers, which is DB2 — but as 60,000 installed sites that might need a product [from another pillar]."

Siegesmund said that allows other IBM products to be leveraged into her group's sites. Another advantage is IBM's new focus on small-to-medium businesses, "which is where a lot of my partners sell," she said. "IBM is thinking, 'Where do we have partners that sell into that space?' Really, most of the partners that sell into that space are partners that came from Informix, including the U2 partners. So there's a real push to help those partners worldwide be successful and go out and sell more."

One of the real ways that IBM helps its partners to be more successful, Siegesmund explained, is by providing the latest technology for partners to incorporate into their products so they can stay on the cutting-edge within the vertical market in which they sell.

Banking on its name and logo, IBM doesn't advertise anything below the pillar name level, Siegesmund said. "It's the partner's solution that's actually being marketed in the end. Our business plan is to really promote and help those partners through delivering the technology and showing them a way they can keep their application modern and competitive in the industry."

Siegesmund also takes the personal approach to supporting business partners. "I do a lot of conference calls and speaking at partners' user conferences and meetings, helping them with their customer base and prospects, and anyone who has questions about the future of the product," she said.

Doing its part to add to the legion of MultiValue users, Siegesmund said IBM is actively recruiting new partners to use U2 "because they can develop things and go to market quickly. And as an embedded, low total-cost-of-ownership, low-maintenance database, it can't be beat."

Raining Data and Microsoft *Join Forces*

Raining Data recently showcased their Pick Data Provider for .NET product together with Microsoft Corporation, during Spectrum's 2004 Conference and Exhibition. "Through our strategic relationship with Microsoft, the entire Microsoft .NET Framework platform of products and services is now available to the Pick community," said Mario Barrenechea, Raining Data's SVP, Worldwide Sales and Marketing. "The combination of Pick Data Provider for .NET and Microsoft's Visual Studio .NET delivers a productive and robust solution for building and deploying applications that doesn't compromise any of the features or functionality of the Pick Data Model.

"Pick Data Provider for .NET allows Pick developers to build client/server Windows, Web, Web Services, PDA, and SmartPhone applications with Visual Studio. NET, the most comprehensive software development environment developed by Microsoft. Our offering enhances Visual Studio .NET by seamlessly integrating it with D3 and IBM's UniVerse and UniData environments.

"Pick Data Provider for .NET is the catalyst that makes the bridge between the Pick Data Model, Visual Studio .NET, and Microsoft .NET Framework possible and working together in unison," Barrenechea explained. "In addition, Visual Studio .NET supports multiple languages, which can now be used to write applications to access Pick databases. More development options and accessibility to a larger pool of development resources can help companies running Pick applications embrace the latest technology on the market — quickly."

As testimonial to the extensibility of Pick Data Provider for .NET, David Carlson, senior project manager for New England Computer Solutions Inc. (NECSI), was on hand in Raining Data's booth and demonstrated how his company has taken its UniVerse-based green-screen order job tracking software for contract engineers to a GUI Web application in a matter of a few months.

rum 2004...

He began looking at Pick Data Provider for .NET in August of last year, Carlson said, and by mid-September had "three or four key screens for the application done and working great. Using Pick Data Provider for .NET allowed us to extend our Pick-based application through Web Services. Raining Data's tools gave us the ability to gain direct access to PICK programs and data."

Delegates listen attentively to the Raining Data presentation

Pick Data Provider for .NET has allowed NECSI to provide more options for its customers. "It gives them the simplest basic GUI screens that customers really require and allows them to do more with their externals," Carlson explained. "Before, users needed Telnet sessions to get into their system. Now they can access information from the Web. Thanks to Pick Data Provider for .NET, we are now supplying three new extensions to the application, that weren't part of the original order job tracking software, including the customer's ability to track open job orders via the Web."

Barrenechea said that Raining Data will "continue in our path of delivering excellent value to our customers and the entire community, with our software infrastructure database and connectivity products. The Pick Data Provider for .NET is a sound reflection of this commitment."

Raining Data and Microsoft Deliver

Keynote Address

Raining Data and Microsoft Corp. co-sponsored the MultiValue breakfast the first morning of the Spectrum Exhibition and jointly delivered a keynote address. "For the last year, Raining Data has been working closely with Microsoft to deliver Pick Data Provider for .NET," said Mario Barrenechea, Raining Data's president. "It's very gratifying to see the energy between our organizations and how we've been able to put together the demonstrations we'll be showcasing in our booth."

Barrenechea spoke about the strategy behind Pick Data Provider for .NET (PickDP.NET) and what it means for PICK developers. "I've been around in this marketplace a long time," he expressed. "I started with Prime INFORMATION and Microdata Reality; I'm a developer at heart. I've always been able to distinguish what is hype and what isn't. This product [PickDP.NET] is available today; it's real for PICK applications. It's about maximizing ROI, not about rewriting existing applications."

Revelation Technologies Rolls Out OpenInsight 7.01

Revelation Technologies rolled out Version 7.01 of OpenInsight (OI), its flagship database management system, which was also made available on its Web site the morning the show opened. Mike Ruane, president of Revelation Technologies, stated, "The reception has been very good. People are happy with what they see. They're happy that we're running on Linux now, and they're really happy with our U2 connectivity."

Ruane remarked that the 7.0 release of OpenInsight had 1,700 changes, and the latest release, 7.01, is a patch with only about 40 cosmetic fixes that were needed.

With Version 7.01, OpenInsight has a more modern look and feel. "We had a look and feel that was 10 years old; we updated that," Ruane said. "We improved our XML handling, which is a constant improvement. People get better at it, so they want more things out of it.

"We also rewrote some of our tools internally. We have a new editor, we have new Windows functionality, working with Windows controls. We have real strong Unicode support, so OI works much better than it did before in non-English character languages, such as Russian, Chinese and German."

In the future, he said, the company wants to ensure that the "open" remains in OpenInsight and that it can talk to other databases and have the connectivity that everyone wants. "Because it is so open, it can talk to these different databases fairly easily, pass data from one system to another, and do some value-add in between."

Revelation Technologies also sees a lot of potential in the Linux market for a product such as OpenInsight, Ruane said, and is looking to penetrate that market further.

Staying the Course

Over the last year, Northgate Information Solutions, headquartered in Hemel Hempstead, England, and developer of the Reality DBMS, has been vigorously pursuing the North American market. At last year's Spectrum show, the company announced that it was reintroducing Reality — already a preeminent database in the U.K., dominating the human resources, local government, police and emergency services arenas — to the North American market.

Tim Holland and Rich Lauer, two early figures in the history of the MultiValue market, were central to the plan, as the agents spearheading the marketing effort. Now it seems that the original strategy is undergoing a major overhaul.

Just before Spectrum 2004, Lauer and Holland opted to pursue other opportunities outside of the MultiValue market, leaving Northgate to examine its own options for the future. Alf Pilgrim, chief technology officer, told *Spectrum* magazine that the company was definitely going to continue pursuing the North American market.

> "We build market share by having good applications, doing what we said we would do, meeting customers' commitments, and adding value to the customer base."
>
> Alf Pilgrim, Chief Technology Officer, Northgate Information Solutions

"It's been so much work, so much investment over the last year," Pilgrim commented. "We're here to stay. We would like to see more business, but we are here to stay. We'll continue supporting the show and the MultiValue market."

Pilgrim said it's a "slow burn" to gain a foothold in a new market, spread the word about the quality of the Reality product, and get more recognition as a serious contender to be considered when moving to a new database environment. Reality has been around as long as the PICK market has existed. Throughout a series of acquisitions, the Reality DBMS was maintained and enhanced until it has become the product it is today: a proven, resilient 24 by 7 database management system.

One difference between Reality and other database management systems is that it bundles many features in one package that some vendors charge separately for. "The only thing we don't bundle is the resilience options, Pilgrim said. "They are really an added value and a unique feature to us."

Reality excels in the resilience department, which makes it attractive to computer users in today's high-demand business environment. "Reality meets open standards so users can plug and play where most shops have a mix of relational as well as MultiValue environments," he said. "Another key aspect is its availability 24 by 7; the systems have got to be up."

Northgate built its market share in the U.K. by word-of-mouth, and the company is hoping to do the same in the U.S. "We build market share by having good applications, doing what we said we would do, meeting customers' commitments, and adding value to the customer base," Pilgrim expressed. "And we've been fairly lucky that once we get a client, we very rarely lose them. We have customers going way back that are still with us ... which is a nice feeling."

ONgroup Brings MultiValue into the World of Oracle

ONgroup showcased ONware, and in collaboration with Oracle Corp., gave a talk at the Spectrum conference demonstrating how MultiValue applications can run directly under Oracle without reengineering. Much of the time, effort and expense that would be required to transition a MultiValue application to Oracle has been eliminated with the introduction of ONgroup's ONware product. Rather than reengineering MultiValue applications to run on relational technology, companies can continue to use their MultiValue applications and use relational tools as well as MultiValue tools.

According to ONgroup, whether you are supporting Web-based applications or need better access to your data, the challenge lies in protecting the years of investment tied up in perfecting your MultiValue business applications. ONware is a MultiValue environment like U2 or D3, but it doesn't have an implicit database system, the company said. Instead it can operate on virtually any database, including U2, D3, Oracle and SQL Server or any combination of supported databases. ONware provides the MultiValue BASIC language compiler, dictionary object compiler and executes PROCs, Paragraphs and Sentences. ONware provides a development and runtime environment just like your current MultiValue environment. The user has the choice of where data is stored and in what format. The ONware run-machine manages the database location, database type, and database structure on a file-by-file basis.

In a conference session, Aalok Muley, of Oracle Corp., talked about Oracle's strategic IT initiatives and the DBMS's advantages such as high availability and security. "I've seen a lot of interest from database customers who are building e-commerce sites," he said. "The focus at Oracle is on Web Services standards compliance and being a standard-based platform. Oracle supports all major operating systems and hardware platforms, and it's interoperable with existing infrastructures."

Oracle's enterprise grid computing enables end users to pool resources, manage application service levels, and ensure a high quality of service as well as automate user and resource provisioning. Advantages include quick identification of bottlenecks and acceleration of problem resolution.

"We want to make sure management of applications is an easy task," Muley commented. "Oracle can help circumvent the problem of difficult to manage Web sites. They're easy to manage if you're managing from one console — and you have that with the Oracle portal. The Oracle portal allows multiple applications to be put on a single console. It's easy to use from an end user's point of view and from the point of view of administration."

There's still more news to report from the MultiValue

nts
2004...

mits

MITS Demonstrates Speed of OLAP/Business Intelligence Tool

Management Information Tools Inc. (MITI), based in Seattle, Wash., performed on-the-spot demonstrations of how the MITS product can take a static report and, in seconds, make it interactive, enabling the user to drill down to an amazing array of details. MITS is a reporting tool for gathering, managing, distributing and analyzing data for improved strategic and tactical decision making.

MITS is easy to use and doesn't require extensive training, allowing users to use standard navigation techniques that lets them focus on researching data quickly. "End users no longer have to go to MIS to get reports," said Larry Christensen of MITS. "The end user customizes their own reports."

MITS utilizes "hypercubes," which are derived from your organization's past and ongoing operational information. With this database, MITS allows your users to easily and instantaneously view activity summaries. "You don't have to burn CPU power because the information is already pre-processed—that's the hypercube," Christensen explained. "The user can then show and manipulate data in many different ways including a graph format."

MITS is available in a GUI front-end, Web, or ASCII interface.

jBASE Supports IBM's iSeries

jBASE International announced that jBASE Release 4 supports the IBM iSeries Computer System, opening up new markets for MultiValue VARs. jBASE runs in OS/400 64-bit native mode and uses DB2 and/or the MultiValue database. jBASE is database independent and can also use Oracle, SQL Server, Progress or Cache on other hardware platforms.

The first jBASE installations on iSeries were in 2003. The eServer iSeries Computer System is the leading computer system sold to mid-size and large-scale financial and manufacturing companies. The iSeries allows the user to consolidate multiple Windows servers inside an iSeries to manage all the servers together in a single system and delivers flexible storage management.

Everything the iSeries has to offer can be integrated with jBASE. The conversion to jBASE 4.1 is the first step and requires the same type of learning required for any other MultiValue migration to a new platform, jBASE said. Once the applications are ported to jBASE, they will run on the eServer iSeries Computer System.

You can download a free evaluation version of jBASE Release 4 at www.jBASE.com.

Via Systems

Via Systems *Introduces* XML Product

VIA SYSTEMS in branching out into new directions. In addition to showing the UniVision MultiValue database management system, Via Duct terminal emulation software, and the WebWizard Web Integration Toolkit, Via Systems introduced Quantum XML Highway, from its parent company, EDP Solutions. Robert Catalano, Via Systems' president, explained: "It's an XML highway server that allows you to do secure document exchange, and it will take documents and purchase orders and transform them into XML documents. You can then send them to another company and it can transform that XML document into a document that fits that company's requirements."

The Quantum XML Highway product line will be available soon, Catalano said. "This new direction is outside of the MultiValue market. We like the fact that we also have a product that can be marketed outside of the MultiValue market as well as integrate within MultiValue."

American Computer Technics to Distribute Visage in North America

American Computer Technics (pronounced "Techniques") Inc., Fuquay-Varina, N.C., announced at the Spectrum 2004 show that it will take over distribution in North America of the Visage product, including sales and support. Visage is a browser-based GUI development and data mining tool. A unique product, Visage is a complete development system that allows users to generate screens, applications, and reports, and perform data analysis.

"Visage allows us to do all the things we want to do in today's technological age: data mining, Web deployment, and rapid application development—the most important part," Patrick Williams, president of American Computer Technics, told *Spectrum* at the show. "Visage does all this with safety, security, and data integrity in mind. It's not a slapdash, sloppy system that leaves you wondering about security. And it's truly based on MultiValue database technology. Visage was built to use advantages that we've all come to know and love in the MultiValue community."

Ease of use is also touted as one of Visage's strong points. "You can get started right out of the box," Williams commented, "showing customers fully functional screen designs with a data mining tool right away. It doesn't take years to learn this technology.

"We can do all the flashy things other non-MultiValue products do without losing the advantages of the MultiValue model. It's a complex yet easy-to-use product, and it's what folks have been looking for in this market."

In their conference session, Williams said attendees were most excited about the data mining technology and how quickly they can build significant, multiple views of data. "The most important part is end users can do this. They don't have to keep coming back to the programmers or technical staff for reports. They just click and there it is. It's almost like a view of the business' history that will portend the future of business trends.

"Visage is a strong new technology that catapults MultiValue users into new millennium technology and allows them to do all the things people can do on the Web with colorful, useful screens. Attendees are really interested because they know they will have to have something like this [to stay competitive]."

For more information, visit www.AmericanComputerTechnics.com.

event of the year. Coverage of Spectrum 2004 will continue with Part 2 in the next issue.

Relational + integration: >> it's a Reality!

We all know that in the world we live in today, integration is key to the success of any IT department's strategy. Organizations must be able to integrate disparate systems and enable the sharing of data among an increasingly demanding audience. It's now taken for granted that any serious commercial database must be "open" and supply such interfaces as ODBC (Open Database Connectivity) and JDBC (the JAVA equivalent).

"Organizations must be able to integrate disparate systems and enable the sharing of data among an increasingly demanding audience."

Reality, from Northgate Information Solutions, has provided these facilities for some years now. This article discusses how Reality has shifted the boundaries once more, giving IT departments around the world an "easy" way to integrate data stored in relational databases from within Reality, using SQL-VIEW. We'll also take a look at some of the considerations required when "normalizing" your MultiValue data for presentation and update through ODBC/JDBC.

SQL-VIEW

With Reality V10.0, it is now possible to look outwards and interrogate or update data stored within an ODBC-compliant database. What does this really mean? Well, until now you would have had to "jump through hoops" or purchase third-party products to expose external relational data to Reality users and applications. Now, one simple command can accomplish the task that has eluded developers and users alike for years.

The TCL verb SQL-VIEW creates a Reality file that provides a view of an existing SQL table (or SQL view) on a foreign ODBC-compliant database. This command does not create a new table. The DICT section resides in the local Reality database; the Data section(s) are "live views" of tables in the foreign database.

"You would have had to 'jump through hoops' or purchase third-party products to expose relational data to Reality users and applications."

Within the "live" view created, each row of data is translated into a Reality item while each column is an attribute in the item and Primary Keys are translated into the item ID. Where a table has a single Primary Key column, the Reality item-id maps to this column and is used directly to identify an SQL row. Where a table has multiple Primary Key columns, the Reality item-id comprises the same number of parts and these are used to identify the matching SQL row. The delimiter used to separate the different parts of the item-id is '\' by default, however, this can be changed to a character of your choosing. Dictionary definitions will also be automatically created for each column specified when creating the "view."

With SQL-VIEW, applications can choose to read and/or write (subject to table permissions on external database) data from within Reality. The foreign database will impose strict control over the type and size of data that may be stored in each column. A Reality application using SQL view files must be aware of the format of the external table data and must keep within these controls. Reality will obey the rules of the foreign database for each of the "columns" specified, thus always ensuring the integrity of the data on the foreign database.

Normalizing Existing Data

We will briefly look at some of the issues you will need to tackle when "normalizing" your data. However, this is a huge subject and impossible to cover completely within this article as it's very much data dependent.

SQL is designed to work with a relational database using data structures designed to reduce data duplication and inconsistent data dependencies. The process of designing an SQL catalog in this way is called "normalization."

Reality and SQL data structures are significantly different. The major difference is that SQL is designed to work efficiently with a normalized relational database, whereas Reality is completely independent

> "One simple command can accomplish the task that has eluded developers and users alike for years."

of type and can be efficient and useful even when the data is non-normalized.

Relational databases store data in flat files (tables), within "rows" and "columns" and are strict in specifying the type of data that can be entered into a column. An SQL column definition defines precisely both the type and size of data that can be entered, whereas a Reality data definition only defines the display width of the data, and is completely flexible with respect to the type. This presents us with some interesting challenges in normalizing our data.

Traditionally, relational databases cannot hold sub-field information and there is no correlation in an SQL catalog with Reality multivalues and subvalues. Hence, when converting a Reality database into an SQL catalog, it may be necessary to restructure the data model in order to handle multivalues, or subvalues.

In order to restructure the data model to handle these elements, you need to make rational decisions based upon the data. The column definitions within Reality describing the normalized data for multivalues and subvalues data can be set up in the following ways to either:

- Use first multivalue or subvalue, only
- Use all multivalues or subvalues, separated by spaces
- Use n'th multivalue or subvalue
- Explode—generate a separate row for each multivalue or subvalue.
- Use the last multivalue or subvalue

Examples of Accessing Multivalued Attributes

The following Reality file item contains both multivalued and non-multivalued attributes.

Reality Item	SQL Column Definitions
ITM01	ID
001 ABC]DEF]GHI	ATTR1
002 A002	ATTR2
003 123]456]789	ATTR3

The SQL column definitions for Attributes 1 and 3 are set up for 'E'xploding values. The Reality file only contains one item. Selecting the table produces the following output:

```
SELECT * FROM T1
```

ID	ATTR1	ATTR2	ATTR3
ITM01	ABC	A002	123
ITM01	DEF	A002	456
ITM01	GHI	A002	789

3 Rows listed

The inclusion of a column with exploding values means that there will be a corresponding row for each multivalue position in the Reality item. In addition, the values in all non-exploding columns are repeated on each row displayed. For example, if

continues on page 46

My previous article covered menus, edit keys, passwords, logging on and off, function keys, and the action bar. In this article, I will cover more of the areas you need to be aware of when you are getting started with SB+, including Keys and their use, Troubleshooting: error messages, Keyboard, and Help levels.

Keys and Their Use

The ESC key takes you back one input step unless you have the CUA Arrow movement in SB+ Control Parameters set to Y (Figure 1). If this parameter is set to Y, the ESC key will not take you back one step in the screen, but will cancel the current record and return you to the key field for the screen you are in. If no key field is present on the screen, the screen will be exited and you will be returned the menu. When you set up a new user in user security, enter N in the Suppress Confirm on Screen Escape (Figure 2) so you will receive the confirmation screen when you press the ESC key (Figure 3).

Figure 1 - CUA Arrow Movement Paramater

Figure 2 - Suppress Confirm on Screen Escape

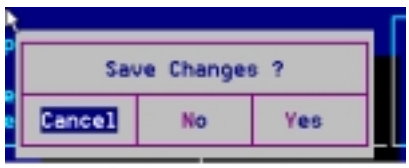

Figure 3 - Escape Save Confirmation

The \ (back slash key) is used to clear a field or make it NULL. The Delete key pressed repeatedly or the Clear to End of Line Key may also be used to NULL a field.

The / (forward slash key) may be pressed if you are at a menu, to request a process name to execute such as MP (Modify Process). At other input prompts you may call a process by entering /processname such as /CALC (Figure 4). If you use the / key on its own at any other prompt it will be taken as its own symbol value.

Generally UPPER CASE characters are required when you enter commands. You should also be aware that System Builder Plus uses F1 through F10 and these function keys should not be used for applications which you develop. If you avoid using these keys you will not be overriding SB+ key assignments.

Figure 4 - / process

◆ F1 will retrieve a short one or two line Help message at the bottom of the screen (Figure 5).

Figure 5 - F1 Help Message

◆ F2 will save the record you are currently working on.

◆ F3 will display intuitive help if it has been defined for the input field you are working in.

◆ F4 will delete the current record you are working on. A confirmation screen will be displayed before the record is deleted (Figure 6).

Figure 6 - Deletion Confirmation

The F1 through F4 keys will work the same in all SB+ processes.

Troubleshooting: Error Messages

Whenever possible, SB+ will validate your input and will display an error message if it detects an invalid input. You may develop input validation routines or use the standard SB+ validation routines such as a list of valid values, a range check or the input exists on another file.

If you get an error message with a number or a message in a different format, it is most likely an operating environment error and is not an SB+ error. Refer to your operating system manual for a description of the error and the action required to correct the error.

Troubleshooting: Keyboard

If your keyboard does not function as you expect it to when you press a key, such as pressing the Up Arrow key and the cursor does not move up on the screen, then you may not be using the correct SB+ terminal definition for your port and the terminal you actually have connected. When you log into SB+ you are asked for the terminal definition you wish to use; make sure to use the correct definition for the terminal you are using. You may not be asked for a terminal definition when you log into the SB+ if the system administrator has set up a default terminal type for your id. If you do not specify a terminal id, the default terminal id DEFAULT.TERM will be used by SB+ (Figure 7). If you enter /KEYS (Figure 8) at any input prompt, you will be able to see what keys SB+ is expecting for the terminal definition you chose when you logged on to the system.

Figure 7 - Default.TERMid

Figure 8 - /KEYS

Help Levels

◆ 1st Reminder. If you press the F1 key a short one or two line Help message will be displayed at the bottom of the screen. If there is additional Help available it will be indicated at the end of the F1 Help message (Figure 9).

Figure 9 - F1 Additional Help

◆ 2nd Window. Pressing F1 after accessing Help will display a window with additional Help, if available.

◆ 3rd Full Page. If there is additional Help available after the first and second level Help is displayed, pressing the F1 key will display a full page of Help (Figure 10). There will be a total page count displayed at the bottom right corner to indicate how many pages of Help are available. You may print the full page Help by entering P and pressing <ENTER> at the bottom of the page.

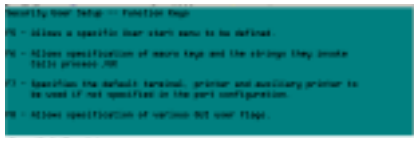

Figure 10 - 3rd Page Help

◆ 4th Help menu. Pressing the F1 key after full page Help will give you a Help menu specific to the current screen (Figure 11).

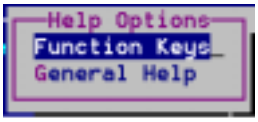

Figure 11 - 4th Help

◆ 5th General Help. There are a number of general topics which are available such as /GLOSS which will give you glossary help.

◆ Intuitive Help. F3 will display intuitive Help at most SB+ prompts, if it is available. When you design an SB+ application it is your responsibility as the designer to provide the intuitive Help for each of your prompts. If there is no F3 Help available, you will receive a message informing you that no Help is available (Figure 12).

Figure 12 - No Intutive Help

When you build applications the facilities and functionality available throughout your application will be dependent on you, the designer. It is incumbent on the designer to build complete, meaningful Help messages for the application. **is**

Danny Passig is a senior software engineer at Natec Systems. He has 33 years experience in the IT field. He has done customer software development and system installations for various businesses. He has worked for IBM as a support engineer for System Builder. Danny holds a Bachelor's degree in Business Administration/Accounting and a MSCIT from Regis University.

SB+ TECH TIPS FROM IBM

How to Control SB+ Report Writer Reports Output Selection Dialog Box to Have HTML Checkbox Not Checked

The default state of the HTML checkbox is set via attribute 2 of HTML.OPTS in DMCONT (0: Not checked, 1: Checked).

If the default in DMCONT, HTML.OPTS is NOT set, 1 is assumed and the default "state" may be set on an individual basis via OUTPUT.SELECT if SB+ is 5.2 or higher.

In SB+ 5.0.5 you can set the default state of the HTML checkbox via attribute 2 of HTML.OPTS in DMCONT (0: Not checked, 1: Checked).

If you have existing reports that you want to control on an individual level, change attribute 37.2 of the report definition (DICT file_name). For older releases, you can change all reports, for example, have a paragraph that selects all of your report definitions, reads each of your report definitions and changes attribute 37.2.

Note: If trying to print a report while in GUI to an F option, it will prompt you to save the report local to the client not the server if the allow HTML flag was not set to N. The allow HTML flag in the report definition F6 is attribute 37.2. This is important if upgrading from older releases and depending on the reports printed to an "F" option to live on the server.

SBClient Tech Alert!

Installing SBClient version 5.3.X on workstations that have any SBClient, version 4.X or older, will make SBClient version 4.X switch into DEMO mode including any SBClient version 4.X installed on the same LAN. If upgrading to version 5.3, we recommend that all workstations using SBClient version 4.X be upgraded at the same time. For further information, please email technical support at u2supp@us.ibm.com.

History of MultiValue

International Spectrum *magazine celebrates* 20 years *of publishing through this new* **"Where Are They Now?"** *series of articles*

Some call the group eclectic. Some say the times were eccentric. No doubt, the early group of PICK dealers who sold Reality on Microdata in the 1970s was a charismatic group of passionate salesmen, technical visionaries and entrepreneurial risk-takers. Their common goal was to build a better mousetrap (minicomputer solution) for small to medium companies. Their common obstacle was the very company that helped set them up, and then tried to squeeze them out. The results? A story of huge opportunity, shortsighted greed, a splash of competitiveness, and long-lasting camaraderie. *Spectrum* magazine takes a look at this original Microdata lot of dealers, where they came from, what they achieved, and what they are doing today.

In celebration of *International Spectrum* magazine's 20th anniversary, this regular "History of MultiValue" column will be focused on notable "characters" and "personalities" throughout the history of the MultiValue marketplace. In this issue, we're taking a look at the industry's original dealer network.

BY LEANNE GREEN

John Keogh Helps Microdata Build Its Dealership

In the early 1970s, Dick Pick entered into an agreement with Microdata to implement his software concepts on Microdata computers. Pick's system was renamed REALITY and became the exclusive property of Microdata. In 1973, Microdata hired John Keogh, one of the original "four" at Basic Four, to head its new Reality on Microdata sales efforts. Keogh had been implementing a dealer channel while at Basic Four, but left before its fruition. He picked up at Microdata where he'd left off at Basic Four, and set out to create one of the first dealer channels of its kind in the computer industry.

Aware of the opportunity at hand, Keogh took a unique stand as a Microdata employee. He not only launched the dealer network by finding and signing the first group of nationwide dealers, but also told Microdata that he wanted the local, Southern California dealership for himself. Microdata actually provided the seed money for Keogh's new dealership.

Keogh spent two years signing up what he recalls to be "the first 10 to 15 original dealers." A fair number of these first dealers had been Basic Four software houses, so Keogh was building what he felt was a strong foundation of known successful business people. "The basic quality I was looking for in a potential Microdata dealer was an understanding of software and how to install a computer," Keogh commented.

Although Keogh had begun a similar program at Basic Four, many believe the Microdata dealer network was the first of

its kind in the computer industry. Keogh admits that he wasn't aware of any dealerships before Microdata. "The whole concept of a dealership is that it allows you to employ entrepreneurial guys who put their own money and sweat into the deal, and as a result, makes you a lot more successful than if you opened up branches with employees," Keogh said.

Keogh proved to be successful in finding the right group with plenty of money and sweat.

"The original PICK dealers were strong, aggressive businessmen that sold a lot of product," remembered Rich Lauer, who helped Keogh run the Southern California Microdata dealership. "For one reason or another, John was successful and lucky in finding the original guys and putting Microdata on the map."

Ted Sabarese was one of the first to buy into a Microdata dealership from Keogh. "The truth of the matter is, the original dealers were all really smart guys," Lauer said. "Keogh picked them. And Keogh was a smart guy."

Great Hardware Margins + Great Technology = *Great Opportunity*

ASK ANY OF THESE ORIGINAL DEALERS what opportunity they saw in Microdata and Reality, and most admit it was the PICK technology and the hardware margins that initially drew them in.

"Companies like Burroughs were touting the BC6 and BC7 that would run two jobs simultaneously, while IBM was also talking about its new foreground and background capabilities. Microdata, on the other hand, could run 32 jobs simultaneously," said Bob Jordan, an original Microdata dealer.

George Ridgway was one of the first Microdata dealers on the scene in 1973: "Microdata gave us a chance to control our destiny, gave us hardware margins, gave us a wonderful data management system, and at that time, an operating system, and didn't charge us any application fees like Basic Four," he said.

A typical Microdata system would sell for $100,000 to $200,000, with dealer margins falling between 30 to 40 percent. Ted Sabarese said: "Those days are gone—selling hundreds of thousands of dollars worth of hardware and making 45 percent, then selling software up to $100,000 that was theoretically all profit. It was fun. The guys were fun. Everything about it was fun."

Sabarese said the opportunities were prolific because the market was ready. "We weren't all so smart, but the truth of the matter is, the timing for the product was right. There were no PCs. We were selling to companies that either wanted off their mainframes, or had never computerized. When a dealer was selling to an end user, he wasn't giving a dissertation on computer equipment. He was really selling a solution. And that didn't exist at that time. On top of that, we were selling a minicomputer for about $100,000. Previously, businesses had been faced with going to IBM and buying a mainframe for millions of dollars."

Bob Jordan said it became easy for the early dealers to "make quota above quota" by the add-on business. "In those days, if you just stayed in business, you'd get about 30 to 40 percent add-on business for just being alive. Your customers just kept calling up and buying more stuff."

The road to profitability, however, did take some time, technical ingenuity, creativity, and sometimes, hocus-pocus.

Judd Van Dervort remembers his dealer peers as having strong sales backgrounds, a lot of charisma and entrepreneurial spirit—important traits when you were selling the Microdata product into businesses without any business software to sell. "They [Microdata dealers] were good at making something out of nothing, wheeling and dealing, and getting a lot of business," he said.

This lack of software made for a rocky road in the beginning. "We would often lose money developing solutions because we just didn't have any software," Van Dervort said. "I got into the habit early on of just demonstrating the database and convincing the client that they could develop their own software."

"In those early days, there wasn't a lot of software around, so we were making it up as we went along," Lauer confirmed. Sabarese agreed: "Everything had to start from scratch. There was no vertical niche. We were general software houses and were just writing custom programs."

An Alliance of Solidarity and Camaraderie Among Dealers

After the first few years operating under the Microdata dealership network, the dealers formed the Microdata Dealer Association. It was the mid-1970s and the dealers were beginning to have trouble with Microdata. The association met a few times a year without Microdata and became a place for sharing ideas, resolutions, friendship and a few drinks.

Ted Sabarese, who started the Microdata Dealer Association with the help of John Keogh, said the original group of dealers always worked very well together: "[All the people named in this article] are really good guys, we liked each other, we'd socialize at the dealer meetings, and if problems came up, we'd call each other and talk about it. And Dick Pick also liked this original group of dealers. He even paid the dealer fee just so he could come to the dealer meetings!"

Tom Davidson said the dealers were a good group of guys, fighting a losing battle with Microdata. "It would have been far more difficult for any of us to have survived the conflicts and poor management between Pick and Microdata, had we not been able to have those meetings and share our frustrations and come up with solutions that could work."

Dave Rodgers agreed with Davidson. "We had a lot of fun together as dealers and really enjoyed each other's company. And as it got more and more combative with Microdata, it became apparent that it was easier to deal with them as a group, than individually."

The Demise of the Microdata Dealerships

The first Microdata dealerships began in the early 1970s and after just five years, came to an acrimonious end. Hardware sales were strong, the database was becoming more solid, and the dealers had begun to carve vertical niche markets in growing business segments. This wasn't good enough for Microdata, as it began to change the rules, compete with the dealers, and try a variety of tactics to rid itself of the original dealer network.

Rich Lauer tried to explain Microdata's motives: "When Microdata was acquired by McDonnell Douglas in 1978, it wanted to sell REALITY from branch offices using a direct sales force, instead of through dealers as it had been doing. In some cases Microdata went ahead and in whatever fashion it took, stopped doing business with the dealers."

Many of the original dealers, who still clearly recall the sequence of events that led to the extermination of Microdata's dealerships, paint a clear picture of a company whose greed got in the way of business:

Randy Jordan: "Microdata decided they wanted the territory for itself, so it tried to kill off the dealer organization. And the way it did this was to split us into sub dealers, so it could take back the lucrative territories."

Randy and Bob Jordan share their passion for boxing. Pictured here: Bob is on the far right in the blue and red jacket; Randy is next to him wearing the business suit.

Judd Van Dervort: "McDonnell Douglas was trying to take back some of the larger territories. It tried it by throwing the dealerships out when they weren't performing to their quota. So, knowing that Philadelphia would be a territory it wanted, I met with Microdata and said, 'Why fight me? Why not just buy my dealership."

Judd Van Dervort puts as much time into his 10 grandchildren as he always has with Keystone.

Ted Sabarese: "Microdata drove the dealers away: its policies, its pricing, the ways it tried to go around the dealers. It wouldn't protect the dealers' integrity. Microdata would give you the territory, but if a juicy deal came down, they would try to get it directly. It did all of the 'no-no's you shouldn't do if you're working with independent dealers."

Ted Sabarese became the offshore national champion in boat racing in the years following his championship with Ultimate.

Tom Davidson: "Microdata was very poorly managed from the top down. The company did unethical things, like introduce a new processor that runs PICK, but tell us we're exclusive only on the old machine. Of course, they'd fail with the new processor, but that's where they had put all their money. It became apparent that in the long run, Microdata was not going to succeed and we had no future with PICK on Microdata equipment. That's when six of us from the original dealer group got together and started what ultimately became VMark."

Tom Davidson in his SCCA Spec racer at the Sears Point Raceway in Sonoma.

A Story Worth Telling

The truths behind the original Microdata dealerships; the accounts of the men and women who took a risk on a new technology and distribution channel; and the answers to the question, "Where are they now?" make a great story for the MultiValue marketplace. The fact is, this industry made a lot of people wealthy. And it continues to do so. So our suspicions have been confirmed—companies that find a good, solid niche in a vertical market are often guaranteed the pot of gold at the end of the rainbow.

Although we didn't get to speak with all of the original Microdata dealers (see the sidebar on page 19), we've heard that many more of these entrepreneurs are looking back on their careers with a smile, sipping a cool drink from their vacation home on the Riviera.

NEXT ISSUE? ***Find out whatever happened to the "Most Influential People in PICK" who graced the cover of Spectrum magazine 15 years ago. Once again, our investigative reporters will be back at work, asking the question, "Where are they now?"***

The Dirty Dozen

Spectrum magazine was able to "piece together" this list of 12 "original" Microdata dealers. More of these dealers exist, however, and we'd be glad to hear from you if you remember names or stories that you'd like to send our way.

Microdata Dealer's Name and Original Geographical Territory

- **Joe Apprendi*
- *Don Breidenbach* San Francisco
- *Tom Davidson* Washington metropolitan area, North and South Carolina
- **Wally Haugaard* Pacific Northwest
- *John Keogh* Southern California
- **Randy Naylor* Boston
- *Lou Pizzigone* New York
- *George Ridgway* Illinois, Northwest Indiana, Iowa, Missouri, Minnesota, Wisconsin, North and South Dakota, Eastern Kansas, Eastern Nebraska
- *Dave Rodgers* Raleigh, North Carolina
- *Ted Sabarese* New Jersey, Atlanta
- *Donn & **Sharon Ulmer*** Michigan
- *Judd Van Dervort* Philadelphia
- **Bill Walsh* Colorado

Names are listed alphabetically by last name.
*Those with an * by their name could not be reached.*

John Keogh

Original Dealership Name: Southern California Data
Original Dealer Territory: Southern California

Dealership Evolution:

◆ Founded the company in the 1973-74 timeframe as one of the first Microdata dealerships.

◆ Within five years it grew to be a $7 million per year business with 25 employees.

◆ Dissolved the company in 1979 when he co-founded The Ultimate Corp. and sold his sales territory back to Microdata for about $1 million.

Interesting Facts:

◆ Because Keogh was busy wooing new dealers across the country, he hired one of his previous Basic Four employees, Rich Lauer, to run the new dealership for him. Lauer was the first employee of Southern California Data.

◆ He recalls establishing 10 to 15 of the first Microdata dealers, but left Microdata as an employee in 1975 to focus on running his own dealership.

◆ Co-founded The Ultimate Corporation in 1979 with Ted Sabarese; put the company on the New York Stock Exchange in 1981.

◆ Resigned from Ultimate in 1987 a very wealthy, retired man.

Quotable Quote:

◆ "When Ultimate went public, Ted and I were the two largest shareholders. I've lived off of that since that time."

Where is he now?

◆ Has a house in Mexico; North Scottsdale, Ariz.; and Flagstaff, Ariz.; where he splits his time depending on the weather conditions.

◆ Dabbles in real estate developments and investments.

◆ Bought a company (a manufacturer of ergonomic accessories) with his younger son four years ago, now run by his son and daughter-in-law.

George Ridgway

Original Dealership Name: Systems Management Incorporated (SMI)

Dealership Territory:

Started with Illinois and Northwest Indiana, then took on Iowa, Missouri, Minnesota, Wisconsin, North and South Dakota, Eastern Kansas, Eastern Nebraska

Dealership Evolution:

◆ Began as a Microdata dealer in 1973; stayed with Microdata exclusively until Microdata dissolved the dealerships around 1978.

◆ Ridgway founded his company in 1969 as an applications software development and contract programming company specializing in Cobol, and later, BASIC programming for Basic Four.

Interesting Facts:

◆ Ridgway originally turned Keogh's invitation down, and then spent more time looking at the Reality system through the eyes of his technical vice president, Tim Holland. Ridgway recalls that Holland told him: "This system seems powerful and supportable, with faster application development capabilities. You'd better not turn this down too fast."

◆ When Pick licensed the product to AT&T (PickTel), the company only sold about 12 machines before it folded. Eleven of those were sold to SMI.

◆ SMI developed the RPL programming language for PICK (Real-time Processor Language).

Quotable Quote:

◆ "Pick and Microdata gave us the technology at a reasonable price. It's really what we did or did not do with that technology that made the difference. It wasn't Pick's responsibility to get us sales, it was Pick's responsibility to produce, enhance and support the finest data management system in the industry."

Where is he now?

◆ Still based out of Des Plaines, Ill.

◆ Changed company name to Real Time Software in 1993 when SMI merged with two other companies for financial strength; has 10 employees.

◆ Still involved as a turnkey dealer and software publisher, focuses exclusively on D3 from Raining Data.

Ted Sabarase

Original Dealership Name: Minicomputer Sales and Leasing (MSL)

Original Dealer Territory: Started with New Jersey, then acquired Atlanta

Dealership Evolution:

◆ Began as a Microdata dealer in 1973.

◆ Prior to MSL, Sabarese ran Diversified Computer Technology, a software house developing solutions for Basic Four and Singer's minicomputer.

◆ In 1978, among financial disputes with Microdata, MSL was "assimilated" back into Microdata, at the same time Sabarese was founding The Ultimate Corp.

Interesting Facts:

◆ Retired from Ultimate in 1989 at the young age of 49.

◆ Boat racing became a passion for Sabarese during most of the '90s: He was the offshore national champion for two years.

Quotable Quote:

◆ "The Microdata days were a great time in my life. It was fun, challenging, I met a lot of great people, and built a successful company from that. It's a fantastic success story: We went from a $2,500 investment into a quarter of a billion dollar company. That's all part of the PICK story."

Judd Van Dervort, Sr.

Original Dealership Name: Keystone Data Systems | Original Dealership Territory: Philadelphia

Dealership Evolution:

◆ Signed on as a Microdata dealer in 1974 while working with a professional services company called Keystone Computer Associates.

◆ When his company wanted out of the "fledgling" Microdata dealership, Van Dervort found some venture capitalists to put up seed money, and in 1975, founded Keystone Data Systems.

◆ In 1980 he sold his exclusive marketing rights back to Microdata, used the money to buy out his partners, signed on as the first reseller of Prime INFORMATION, and renamed the company Keystone Information Systems, as it has remained for 24 years.

◆ Keystone, now based on UniVerse and UniData, has been successful by focusing on the vertical markets of school districts (K-12), and local government (public safety).

◆ Today Keystone has 50 employees, with offices in Raleigh, N.C., and headquarters in Southern New Jersey.

Interesting Facts:

◆ In his quest to put PICK onto UNIX, Van Dervort co-founded VMark Software with five other Microdata dealers (see Tom Davidson).

◆ In 1998, Keystone acquired the mid-range local government business unit of EDS, which had formerly been a competitor, Infocel Inc. (see Dave Rodgers).

Quotable Quote:

◆ "U2 has a long future with IBM because it's still selling well and makes IBM a fair amount of money. The 'PICK' world, in its current manifestation, has a lot of legs. It's going to be around for a long time."

Where is he now?

◆ He and his wife live in Moorestown, N.J., a bedroom community of Philadelphia, Pa. They have a home on the eastern shore of Maryland, on the Chesapeake Bay.

◆ Van Dervort is still active in the business. He has, however, turned over much of the day-to-day operations to Keystone's management team, led by his son, Judd Jr. According to Van Dervort: "Judd Jr. will become the president in the next few years, but I will continue to be very involved in the company, which I love, even after I relinquish the role of president and become chairman of the board."

◆ Now that he has more leisure hours, Van Dervort and his wife, Mary, enjoy travel, golf, and "mixing it up with our 10 grandchildren." ■

Where is he now?

◆ Moved to Florida in 1990 where he ran a friend's electronics' firm for three years.

◆ Today he owns an air conditioning company (with his brother) and five restaurants (with a partner) in Ft. Lauderdale, Fla.

Tom Davidson

Original Dealership Name: Datatel Inc.

Original Dealer Territory: Washington metropolitan area, North and South Carolina

Dealership Evolution:

◆ Signed on as a Microdata dealer in 1973, evolved out of the Microdata business around 1978 in a move to the Prime INFORMATION platform.

◆ Founded Datatel in 1968 as an office products provider, was about a $4 million company before it became a Microdata dealer.

◆ In 1971 Datatel merged with another company owned by Ken Kendrick (Kendrick ran the bank processing side of the business while Davidson ran the Microdata side. Today Kendrick is part-owner of the Arizona Diamondbacks baseball team.)

◆ Datatel built its success on a strong vertical focus in the college and university marketplace.

◆ This year, Datatel revenues are close to $100 million and is the most profitable company in the software industry.

◆ Davidson and Kendrick transferred quite a bit of ownership to Datatel employees (70 percent of the employees own part of the company), but the two of them remain the largest shareholders.

Interesting Facts:

◆ Davidson helped found VMark Software in the late '70s in his quest for a processor-independent system. His investing partners were fellow Microdata dealers Donn and Sharon Ulmer, Randy Naylor, Judd Van Dervort, Bill Walsh and Dave Rodgers. VMark eventually merged with Unidata, became Ardent Software, was bought by Informix Software and is now owned by IBM.

◆ Although Davidson was one of the founders of VMark, his company became the largest dealer for Unidata (VMark's key competitor). At that time, the developers that worked for Davidson felt Unidata was a more stable company than VMark.

Where is he now?

◆ Davidson left Datatel in 1989, but claims that he is "not retired yet with a major personal interest and financial interest in Datatel, the company I founded 35 years ago."

◆ In 1991 he moved to Santa Barbara, Calif., where he joined two partners to form Invision Medical, a company specializing in camera-driven endoscopic surgery products. He sold that company in 1995.

◆ Davidson founded the Balance Bar company after he and partner Barry Sears (author, "Enter the Zone") went separate ways. He took the company public in 1998 and sold the company in 2000 to Kraft Foods (a subsidiary of Phillips Morris) for $264 million.

◆ Davidson claims to be out of the business of building businesses, but now enjoys the hobby of building and selling world-class homes. Recent home sales have exceeded $11 million each.

Continues on page 22

Dave Rodgers

Original Dealership Name: Distributed Data Systems
Original Dealership Territory: Raleigh, N.C.

Dealership Evolution:

◆ Began as a Microdata sub-dealer in the early 1970s under Datatel with partner, Skip Shippey.

◆ Distributed Data became its own Microdata dealership after one year under Datatel.

◆ When Microdata eliminated the geographical territories, Rodgers' dealership began to focus on local government applications.

◆ By 1981, the company had made some business acquisitions and was eventually acquired by Adage in Boston. A year later Adage decided to eliminate the PICK business, and with the aid of venture capitalists, Rodgers was able to buy back the business.

◆ Skip Shippey left when they sold the company to Adage. He started another company called Management Systems Associates (MSA), a successful hospital solutions provider. Today the company is A4 Systems in Raleigh, N.C.

◆ Rodgers renamed the company Infocel in 1983. By this time, Infocel had a successful hold on the local government vertical market.

◆ EDS bought Infocel from Rodgers in 1992 for $22 million.

Interesting Facts:

◆ Rodgers and Shippey began Distributed Data Systems with just a few prospects. "Skip did the programming, I did the selling," Rodgers said. The first year they brought in almost $300,000 in sales. By 1980, sales had escalated to $2,400,000.

◆ Rodgers was one of six partners to invest in and found VMark Software (see Tom Davidson).

◆ With partners Shippey and Bill Walsh, Rodgers helped to form a "buying consortium" called National Information Systems. "We couldn't get good margins on terminals or printers from Microdata or Prime, so we organized NIS just to pass through the orders. We were selling to about 150 various organizations and dealers around the country."

◆ Rodgers owned a computer service bureau before his Microdata dealership, which had already reached about $10 million in annual sales. Prior to that he was with IBM.

Where is he now?

◆ Rodgers stayed with EDS until 1994 to help through the conversion, then retired at the age of 62.

◆ Today he lives in Raleigh, N.C., spends some time in the stock and real estate markets, travels a lot and spends time with his wife at a beach house in Wrightsville Beach, N.C.

Bob & Randy Jordan

Original Dealership Name: Tidelands Data Products

Original Dealership Territory: Houston and its five adjacent counties

Dealership Evolution:

◆ Signed on as a Microdata dealer in 1975. The Jordan brothers had two partners in the initial dealership: Marty Martin and Len Krisman (Len left the business in 1978; Marty passed away in 1980.)

◆ Its first year in business, Tidelands made $400,000 profit, with sales over $1 million; by the second year the company was already over $3 million in sales.

◆ The Jordans continued as a Microdata dealer until 1980, when Microdata began "squeezing out" the successful dealerships, and the Jordans then became Ultimate dealers.

◆ Company name changed to Jordan and Associates in 1980.

◆ Installed its first government application in 1981 and have focused on the county. courthouse vertical market ever since. Clients include court systems, tax offices, sheriff offices, small claims courts, etc.

◆ Renamed the company Government Service Automation in 1991.

Interesting Facts:

◆ In the mid-80s, Jordan and Associates sued Ultimate over contractual issues that became heightened when Sabarese was easing out and Mike O'Donnell was taking over. "Our legal bills were running $50,000 per month," remembers Randy Jordan. "That was a lot for us, but not for Ultimate. They had more money than God!"

◆ The Jordans wrote a database to run PICK code on Tandem hardware, with the intent of setting up their own dealer base. Ultimate bought it for a little less than $10 million, along with the assurance of future royalties. "We had kissed and made up by that time," Randy said. Unfortunately, Ultimate was dying. It gave a few copies of the Tandem database to VMark, but never sold any. Randy explained the rest: "Tandem eventually bought the database from Ultimate, but Compaq bought Tandem and Hewlett-Packard bought Compaq. We never got royalties and have no idea where the product is today!"

◆ The Jordans' database was the first of its kind to be seamlessly integrated into the operating system. Before that, all PICK-based databases were sitting beside the operating system.

◆ Randy attributes Bob's success as a salesman to his experience in selling Bibles door to door in high school and college.

Quotable Quote:

◆ "Microdata killed itself. Instead of letting the dealers work out there and make money, the company got greedy and tried to squeeze out the dealers and take back territories and put dealers in little 'po dunk' places."

Where are they now?

◆ The Jordans have homes in Houston, Texas, and Nashville, Tenn.

◆ Bob and Randy's company, GSA, is now a corporation owned by shareholders. It is based in Nashville with offices throughout Texas. With 50 to 75 employees, the company's revenue is in the "eight digits."

◆ Bob serves as vice president of Sales and Randy is the senior analyst; Mike Boswell is the president.

◆ The Jordans said they have gone through all the big toys, including owning their own jets. Today they've settled on a half-million dollar silver engine bus to indulge customers or prospects in luxurious outings.

◆ Both brothers have a passion for professional boxing. Randy has worked as a promoter. Bob works with professional fighters, including Eric Griffin from the 1992 Olympic team and world champion, Randy Johnson.

Donn & Sharon Ulmer

Original Dealership Name: Comtec

Original Dealership Territory: Michigan

Dealership Evolution:

◆ The dealership was originally incorporated as Century 21 Data Systems in 1973, but renamed Comtec a few years later.

◆ It was founded by husband and wife team Donn and Sharon Ulmer, who eventually added a third partner, Donn's brother Greg Ulmer.

◆ In 1977 Comtec wrote and sold its first HMO (Health Maintenance Organization) application. A year later, two other HMOs bought the software and Comtec found itself as an HMO turnkey software solution provider. Comtec quickly became the managed care industry's market share leader.

◆ Within 13 years, Comtec had grown from two people to a staff of just under 70 people. Its Managed Care base had climbed to 90 installations, including almost every major HMO chain.

◆ In 1986, Comtec was acquired by Computer Science Corporation.

◆ After the Computer Science acquisition, Comtec became CSC-Comtec. A few years later CSC-Comtec became the foundation of the CSC Healthcare division, which still exists today. While CSC does not actively market the product developed under PICK, it still supports a large MultiValue installed base.

Interesting Facts:

◆ The Ulmers were one of six Microdata dealers to invest in and found VMark Software (see Tom Davidson).

Where are they now?

◆ After working for CSC for a few years, Sharon left and formed Strategic Alternatives (SAI), which became a master reseller for Digital Equipment Corporation.

◆ SAI changed its name to Whitehill Technologies, and sells mcPLUS, a managed care software solution, to small PPOs.

◆ Donn Ulmer has exited the technology arena and now develops commercial real estate (retail shopping centers and office buildings). Sharon also runs a small property management firm which leases space in these projects to tenants as well as maintains the centers and buildings. **is**

Non-Procedural Reporting in OpenInsight

Part 2 of 2

You Want That for Here OR TO GO?

BY MIKE RUANE, REVELATION SOFTWARE

There are three different ways that developers can generate reports without writing any code.

Users still love their reports. Since Part 1 of this article came out, hundreds if not thousands of new reports have been created.

Imagine if developers had to spend time writing code to support or create those reports? They wouldn't have time to do any other changes or enhancements to the systems they were supposed to be supporting.

In Part 1 of this article, we covered some of the ways that reports can be created, run, and saved in OpenInsight without writing any code. First, we reviewed OpenInsight's System Monitor, where users or developers can enter in simple list statements. Next, we covered OpenInsight's Report Builder+, a great GUI-driven front-end that creates list-style reports, but with control over paper orientation, shading, font, color, and more. We discussed that this tool can send output to the Printer or Screen, but that it had much more potential. In this, Part 2 of the article, we will discuss how this reporting tool can send output to PDF or HTML. We will also show how to report from OpenInsight directly into Excel using our Web tools.

The OpenInsight Report+ Builder and PDF

As discussed earlier, the OpenInsight Report Builder+ can report to the screen or printer easily. And since we use the Windows print manager, your reports can print on any printer that your workstation knows about. However, users often have the need to create a report in a format that can be stored on disk, emailed, and loaded onto a Web site. For this type of functionality, users often require the use of PDF files.

PDF files, short for Portable Document Format, are often viewed using Adobe® Acrobat Reader®, a free piece of software from Adobe Systems. PDF

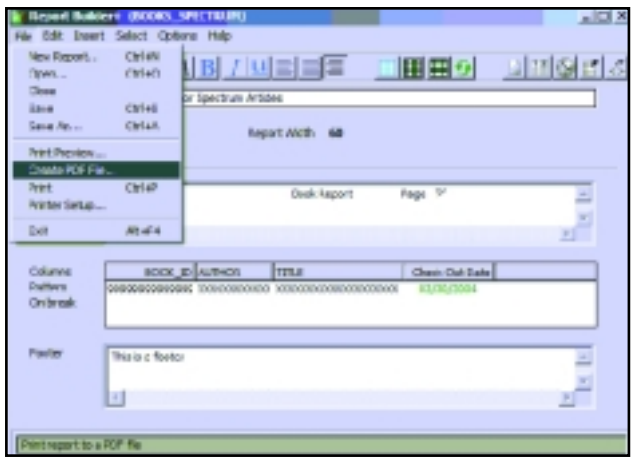

Figure 1: Choosing to create a PDF file at Design Time

files are viewable on most major operating systems, and have become a de facto standard for document interchange. It is used and encouraged for use by business and government alike, all the way up to the United States Supreme Court.

OpenInsight can generate its reports from the Report Builder+ in PDF format. From design mode in the Report Builder+, one can choose the menu choice File-Create PDF File as seen in Figure 1.

Once that menu item has been chosen, the user is prompted for the file name and location where the PDF file should be created, as seen in figure 2.

Once the filename and path have been entered and the OK button clicked, the report previews to the screen and generates the PDF file at the same time. The PDF file is then ready for viewing, emailing, or posting to a Web site.

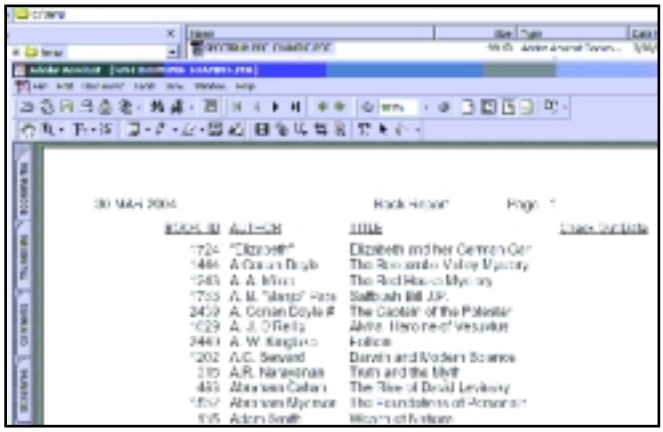

Figure 3: The completed PDF output

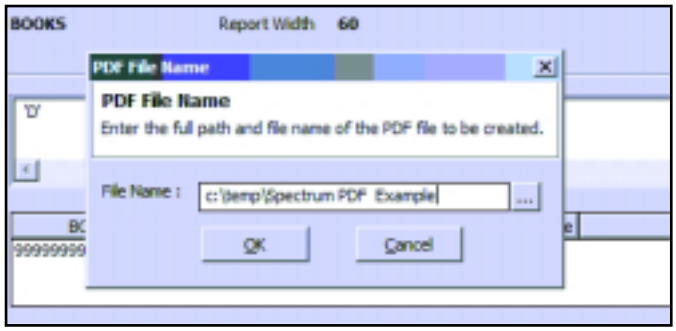

Figure 2: Specifying the output location for the PDF file

Running Reports to a Browser

The Report Builder+ also gives us the opportunity to run directly to a browser. Once a report has been made in the Report Builder+, it can be run in a browser. This is a very powerful feature. Reports that users make for themselves to run while at their desk can be run from remote sites as well, without rewriting them.

Figure 4 shows an HTML code snippet that will launch a report across the Web.

Continues on page 26

Continued from page 25

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 3.2//EN">
<HTML>
<HEAD>
<TITLE>Reports</TITLE>
</HEAD>

<FORM ACTION="http://127.0.0.1/cgi-bin/oecgi.exe/inet_dlist" METHOD="POST">
<INPUT TYPE="submit" VALUE="Blue Spectrum Report">
<INPUT TYPE="hidden" NAME="REPORT_ID" VALUE="BOOKS_SPECTRUM">
</FORM>

</BODY>

</HTML>
```

Figure 4: The HTML used to launch the BOOKS_SPECTRUM report in a browser

When the report is run across the Web, the output appears much as that in Figure 5. At this time the HTML reporting using Report Builder reports does not support colors and shading in fonts; it is expected to be built into OpenInsight Release 7.1.

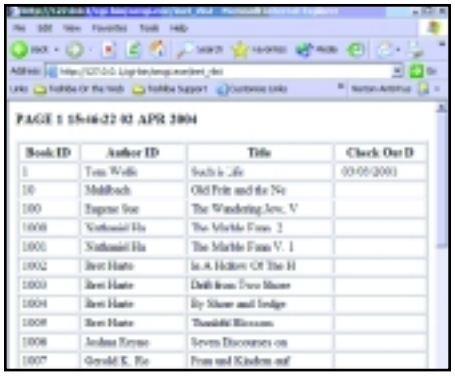

PAGE 1 15:40:22 02 APR 2004

Book ID	Author ID	Title	Check Out D
1	Tom Wolfe	Such is ..life	03/05/2001
10	Muldoch	Old Pete and the Ne	
100	Eugene Sue	The Wandering Jew, V	
1000	Nathaniel Ha	The Marble Faun 2	
1001	Nathaniel Ha	The Marble Faun V. 1	
1002	Bret Harte	In A Hollow Of The H	
1003	Bret Harte	Drift from Two Shore	
1004	Bret Harte	By Shore and Ledge	
1005	Bret Harte	Thankful Blossom	
1006	Joshua Reyno	Seven Discourses on	
1007	Gerald E. Ha	Fran and Kindem and	

Figure 5: Report Builder+ Report run in a browser

Running Reports Directly to Microsoft's Excel

A major request from users is the ability to run their reports into Excel. This may not necessarily be the best practice from a business point of view, as once in Excel, data can be modified. This does not necessarily comply with the Sarbanes-Oxley Act, although it may have some implications. Nonetheless, users want to report into Excel.

Using a Comma Separated Value (CSV) export, one can get data fairly easily, but wouldn't it be better if a user could open up an Excel spreadsheet, either initiate a canned report, or type in their own reporting command? We believe so, and have an answer to it.

Since Office 97, Excel has had the capability of requesting data via an HTTP request over the Web. This process is called a Web Query, and is really nothing more than a call across the Web that displays data in an Excel spreadsheet correctly. Although Excel has its own tools for creating a Web Query, it is probably easier for MultiValue developers to create the Web Query files from scratch.

Web Query files are really just text files that end with an IQY extension. They are normally stored in the C:\Program Files\Microsoft Office\Office\Queries subdirectory. They can be created using the text editor of your choice.

IQY files typically have four lines, separated by carriage returns. They are:

Line 1 is the type of query (WEB).

Line 2 is the version of Query (1).

Line 3 is the location of the Web document (an INET procedure).

Line 4 contains the POST parameters being passed to the Web document.

This typical Web Query file, called CANNED.IQY, will contain four lines, and look like this:

WEB

1

http://localhost/cgi-bin/oecgi.exe/inet_Weblist

cmd='list books id_sym author title id-supp'

Line 1 is always the word WEB, in capital letters.

Line 2 is the number 1, the version of Web Query supported by Microsoft.

Line 3 is the Web address and OpenInsight Internet (INET) procedure you will call, e.g., INET_WEBLIST.

Line 4 is the specific parameters and values to be passed to the INET procedure.

The INET procedures are simply functions in OpenInsight written in Basic. They have to follow three rules:

1- They must start with the prefix INET.

2- They must accept at least one parameter named REQUEST.

3- It is a function, and must return a properly formatted string of HTML.

There are a number of pre-built INET procedures supplied with OpenInsight, and developers are able to write their own. For the purposes of this article, we will limit our use of INET articles to INET_FORMREAD, which reads data from an OpenInsight file into a Web page, INET_FORMWRITE, which writes data from a Web page into an OpenInsight file, and INET_WEBLIST, which creates columnar reports across the Web using LIST statements, and INET_RLIST, which can run Report Builder+ reports (or ad hoc reports) into a browser. We also presume that the reader has read the article "The Revelation Spin on the Web" by Sean FitzSimons which appeared in this magazine in 2003.

And even though this article is about running these reports across the Web of a LAN using OpenInsight, any MultiValued tool with a Web interface should be able to do the same.

To run the Web query, simply open a blank Excel spreadsheet and choose—Data—Import External Data—Import Data. (The menu choices may be different depending upon the version of Excel being used.) A dialog box will appear as seen in Figure 6.

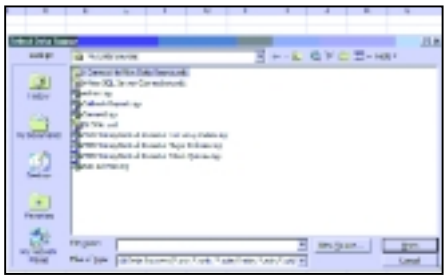

Figure 6: Choose the Web Query to run

In our case, we'll choose the Canned query. Once we have, we are prompted for the location that the data should be put, as seen in Figure 7. We'll accept the default of putting it in Row 1, Column 1.

Figure 7: Choose the location that the data should appear

Once you have chosen the location, the query runs and the data will appear on the screen, as seen in Figure 8.

Figure 8: The results of the Web query

In this example, we listed a field named ID_SYM. You'll notice it in column A, and you will notice that it is a hyperlink. This field is a symbolic, one of OpenInsight's calculated columns, that uses common HTML tags and existing INET functions that are delivered with OpenInsight.

The formula for ID_SYM is as follows:

```
@ANS = '<A HREF = "http://127.0.0.1/cgi-bin/
oecgi.exe/INET_FORMREAD?FORM_ID=BOOKS_
HTM&BOOK_ID=':@ID:'">':@ID:'</A>'
```

(The above should all be on one line.)

You can see it is a symbolic because it evaluates into the variable @ANS.

The pieces are as follows:

<A HREF =	The start of an HTML Anchor Tag
http://127.00.1/cgi-bin/oecgi.exe/	The Address where the OECGI.EXE file resides
INET_FORMREAD	An INET function supplied with OI, that needs 2 parameters
FORM_ID=BOOKS_HTM	The first of two parameters required. A FORM_ID, and its value
&BOOK_ID=':@ID'	Value of the key field for the form named in FORM_ID
'">':@ID:'</A>'	The end of the anchor tag, the key to the record, and the close of the anchor tag. This will cause the value of @id to appear as a link in a browser.

Here's another example showing a symbolic with a report statement embedded in it:

```
@ANS = '<A HREF = "HTTP://127.0.0.1/cgi-
bin/oicgi.exe/INET_Weblist?cmd=LIST
BOOKS ID_SYM TITLE AUTHOR WITH
AUTHOR=':CHAR(39):{AUTHOR}:
CHAR(39) :' ID-SUPP"">' :{AUTHOR}:'</A>'
```

(The above should all be on one line.)

This symbolic, when displayed in a browser and clicked on by a user, will run a report to the browser, listing all books with the author matching the author's name being clicked on.

Finally — Enough reporting options?

So, in this two-part article, we demonstrated three different ways that developers or users can generate reports without writing any code. First, we showed using OpenInsight's System Monitor to enter simple List-type reports. Next we showed the Report Builder +'s capabilities regarding fonts, colors, shading, as well as its output options to the screen, printer, and PDF. We also showed how reports from the Report Builder+ can be run across the Web. Finally, we showed how using Excel and OpenInsight, reports can be run directly over the Web or LAN into Excel. All this without programming.

Now, if you want even better reports, you'll need to cut some BASIC+ code, but that's best left for another article.... **is**

Don't Let SOA Sack Your Business:

Complying With the Sarbanes-Oxley Act of 2002

BY SUSAN JOSLYN, PRESIDENT, SJ+ ASSOCIATES

A man is getting into the shower just as his wife is finishing up her own shower. The doorbell rings and the wife hurriedly wraps herself up in a towel and goes to answer it. It's the next door neighbor. Before she says a word, the neighbor says, "I'll give you $800 to drop that towel that you have on." Naturally the woman drops her towel and stands naked in front of the neighbor, who then hands her $800 and leaves. When she gets back to the bathroom, her husband asks from the shower, "Who was that?" "It was the nextdoor neighbor," she replies. "Great!" the husband says, "Did he say anything about the $800 he owes me?"

***Moral of the story*: If you share critical information pertaining to credit and risk with your shareholders in time, you may be in a position to prevent avoidable exposure.**

Even if the Sarbanes-Oxley Act of 2002 weren't lingering in the shadows ready to put the squeaky clean couple in jail for the miscommunication!

> Senator Sarbanes contends that our SEC legislation — now 70 years old — was hopelessly antiquated for today's business environment.

Named for its main architects, Senator Paul Sarbanes and Representative Michael Oxley, the Sarbanes-Oxley Act of 2002 (SOA) followed a series of very high profile scandals (e.g., Enron), and is, according to President Bush, intended to "deter and punish corporate and accounting fraud and corruption, ensure justice for wrongdoers, and protect the interests of workers and shareholders." Senator Sarbanes contends that our SEC legislation — now 70 years old — was hopelessly antiquated for today's business environment.

What that translates to, for American business, is specific compliance and records-keeping requirements to validate financial records. With a little new incentive: mandatory jail sentences for officers of a company that fails to comply!

Former HealthSouth Corp. chief executive Richard Scrushy was the first chief executive to be accused of violating the SOA. He was indicted last November of overseeing a scheme to deliberately inflate HealthSouth's earnings and assets by more than $2.5 billion over several years. Scrushy has repeatedly denied he had any part in the fraud that left the company battling to avoid bankruptcy, and blamed it on his underlings.

For the last several decades, the American stock market has led the world in terms of volume, liquidity, and scale. An important element of this success has been the credibility of the financial statements companies prepare in accordance to GAAP (Generally Accepted Accounting Principals) for filing with the SEC. Even though the American financial system is arguably still the best in the world, it has not proven sufficient to prevent multi-billion dollar debacles such as those that occurred at Enron, WorldCom, and Adelphia. Individuals at these companies secretly manipulated their general ledgers and artificially inflated their revenues, but their external auditors approved their SEC filings as business as usual. Were the perpetrators clever enough to hide their actions, or did their auditors choose to look the other way to avoid jeopardizing the large consulting revenues they enjoyed from these same clients? When executives at these companies were caught red-handed they claimed ignorance, declaring that they relied on lower-level managers to prepare and present accurate numbers. If these executives were found guilty, their employers could pony up to the bar and pay their fines for them. In hindsight, our system made it all too easy for corporate fraud to occur.

Most of the SOA is focused directly on auditing, responsibility, fraud detection and whistleblower protection. These things seem narrowly focused on legal and financial issues but buried among them are a few little tidbits that open up a whole range of requirements — and opportunities — for IT.

Section 404 of the act — Establishment of Internal Controls — requires companies to recognize, in their annual reports, their responsibility to establish and maintain an adequate internal control structure and set of procedures to ensure accurate financial accounting. The annual report must also contain management's assessment of the effectiveness of the controls they have chosen. Companies are encouraged to deploy both preventative controls (i.e., pre-authorization before funds can be expended) and detective metrics (i.e., reports that summarize after the fact what was spent). All publicly traded American companies will be required to comply with section 404 for their annual reports of years ending on and after June 15, 2004 (deadlines for some organizations were recently pushed back.)

Understandably, CFO's across America rank compliance as their top business challenge and U.S. companies are expected to spend more than $2.5 billion on this compliance in the coming year. This is good news for the lately shrinking IT economy — a significant chunk of that spending is going to information technology projects. The view that this is a finance and not a systems issue ignores the fact that IT systems generate, support, house and transport the financial information, the accuracy for which CEO's are now personally accountable. The CIO must build the controls that will ensure that the data stands up to audit scrutiny.

Since much of our "MultiValue" IT world is in smaller or midrange businesses, less likely to be publicly traded, many of you dear readers are sighing with relief and thinking, "Whew, good thing it doesn't

continues on page 30

Don't Let SOA Sack Your Business *Continued from page 29*

apply to me!" Guess what? It applies to you. It might take a bit longer for it to have a hard effect on you, but SOA compliance is for public companies, companies that may become public (through IPO or mergers) and companies that do business (or wish to do business) with publicly traded companies.

Happily enough, the practices and procedures inherent in this compliance are the same as those which make an organization most efficient. Adjusting for an initial period, of course. Early adopters will have a competitive edge on many levels.

So, what are we talking about exactly? Section 302 of the Act requires that the CEO and the CFO certify quarterly and annually that:

◆ They reviewed the report being filed;

◆ Based on their knowledge, the report does not contain any untrue statements or omit any material facts necessary to make the statements misleading;

◆ Based on their knowledge, the financial statements fairly present the financial position, results of operation and cash flows;

◆ They are responsible for and have created, established and maintained disclosure controls and procedures;

◆ They have evaluated and reported on the effectiveness of those controls and procedures;

◆ Any material weaknesses or deficiencies in internal controls or fraud have been disclosed to the audit committee and the independent auditor; and

◆ Any significant changes in internal controls that could significantly affect internal controls have been disclosed, including corrective actions with regard to material weakness or deficiencies.

Section 408 requires the SEC to review disclosures made by each company at least every three years. A company should be prepared to demonstrate the actions it has taken to ensure compliance and answer potential questions about whether any disclosures are missing or incomplete.

What this boils down to for practical intents is that the CFO and CIO will need to carefully review the internal processes ... all the way down the line. Here's a specific example: Many companies obtain various financial records from various computer systems, both internally and externally. Ultimately a person at a desk uses a combination of automation and manual manipulation to create spreadsheets of consolidated information. Sound familiar? Even 47 percent of public companies still rely on this manual-spreadsheet methodology ... and 89 percent of CFOs interviewed said they had little or no confidence that this type of reporting provided adequate control.

> Even though the American financial system is arguably still the best in the world, it has not proven sufficient to prevent multi-billion dollar debacles such as those that occurred at Enron, WorldCom, and Adelphia.

If you start with that spreadsheet and start working backward to verify the integrity of the data coming into it, you will probably find your nose poked into every aspect of your company's business processes and every computer system known to man along the way. You will find yourself working to eliminate and automate steps, to minimize the possibility of human error (intentional and unintentional!) This is a perfect example of the mushrooming scope of this compliance. From one simple Excel report you may well find yourself investigating and possibly reinventing business processes at every level of your organization.

There are not a lot of specifics to this Act of Congress — for the how's and what's we're going to have to rely on existing frameworks — ISO, IEEE, DOD and CMM to name just a few. But being the clever and resourceful IT professional that you are, you will figure that if you're going to spend that much time and effort on this, the benefits of "world class" process improvement are beckoning. What does it mean to be "world class"? A phrase bandied about by many, its most authoritative definition comes from The Hacket Group. The Hacket Group is the world's recognized leader in best practice research. They have profiled nearly 2,000 organizations, including 100 percent of the Dow Jones Industrials, 90 percent of the Fortune 100, and 84 percent of the Dow Jones Global Titans Index. In their definition, "World Class" means "the ability to structure intelligent business processes that make the best use of technology" and they state that this "... will increasingly be a distinguishing feature of world-class finance." Through their research, they have determined that a "world-class" company scores in the top 25 percent both in terms of efficiency and effectiveness when benchmarked against other organizations. World-class companies in their study had these things in common:

◆ Lower costs: lower finance costs as a percentage of revenue than average companies (.76 and 1.07 respectively).

◆ Higher productivity: if a company designs business processes and reduces unnecessary hand-offs, eliminates manual steps and leverages technology, its people will be able to work faster and smarter. A world-class company has 51 percent fewer full-time equivalents for transaction processing than the average company.

◆ Has eliminated redundant or inefficient control activities.

◆ Has standardized and simplified systems. e.g., limits the number of systems and application interfaces, reducing the likelihood of missing data and time-consuming reconciliation.

Whether you are willing or able to achieve "World Class" status, the changes required to comply with the SOA can benefit the organization in many ways. The ramifications of this congressional act are going to be far-reaching and as the definitions become more and more precise it is going to change the way America and the world does business. The information revolution has run rampant and unchecked, but there is a new world order. "Information Democracy" and "Corporate Transparency" are the rule of the new day. **is**

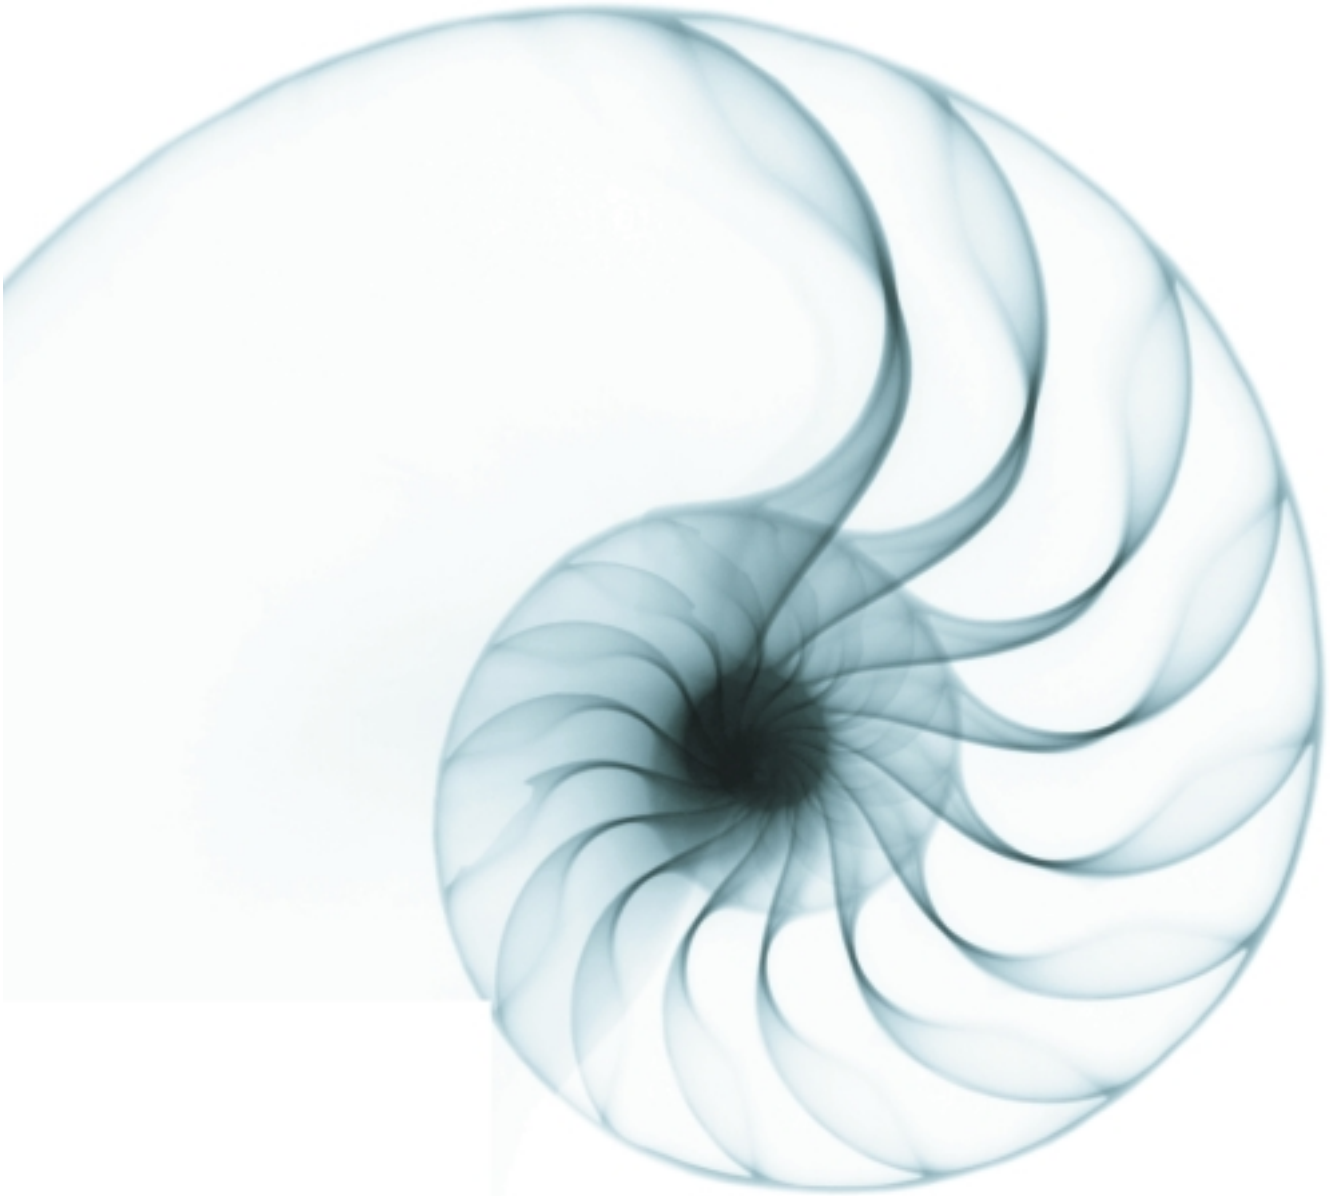

YES

JOIN THE REVOLUTION

"DesignBais provides Evolve 360 with an ability to modernise the user interface of our distribution, supply chain and warehouse management application without significant costs. It is very intuitive so we will be able to secure our investment in existing technologies and leverage the knowledge of our development and consulting skills immediately."

Alan Godby - Product Director - Evolve 360 Pty Ltd

DesignBais

www.designbais.com

newsmakers

jBASE Gateway for Cache Allows jBASE and Cache to Coexist

jBASE International has extended its offerings with the latest addition to the jBASE product family: jBASE Gateway for Caché. jBASE Gateway for Caché is a set of software components that allows seamless access to data held in a Caché database. It provides the ability to access Caché files, in real time, from a jBASE application as if they were native jBASE files — allowing jBASE and Caché to coexist and cooperate.

The unique jBASE jEDI architecture is the enabling interface utilized to interoperate with existing data located in a Caché database. By means of specific jEDI drivers, jBASE BASIC I/O statements can access Caché databases, as simply as they can the jBASE database. Any application written using any tool of the jBASE product suite, e.g., jBASIC, Java or COM OBjEX, can access other data stores as if they were native jBASE files with no change of application code.

With jBASE Gateway for Caché now installed and in operation at a jBASE customer site in the U.K., the list of databases which can seamlessly interface with jBASE applications via the jEDI now includes DB2, SQL Server, Oracle, Progress and Caché.

jBASE has also retained extensive Caché experience and is offering full-service implementations with this new jEDI.

"With the addition of jBASE Gateway for Caché, jBASE International further complements our product set and demonstrates our commitment to offering customers complete solutions," states Pete Loveless, CEO of jBASE International Inc. "We see this as a great opportunity for our reseller channel, opening up new opportunities in both the ever expanding MultiValue marketplace and beyond." Loveless continues, "We will continue to develop and adopt state-of-the-art products that enable increased connectivity and that meet and expectantly surpass our customers' requirements."

jBASE International is now promoting, selling and supporting jBASE Gateway for Caché. For pricing or additional information, please email sales@jBASE.com

About jBASE International

jBASE International is a leading supplier of database management software and Web-enabling tools for developing, deploying, and maintaining business applications solutions. The flagship product, jBASE, was designed from the ground up to be an open database product that would bring the strengths of MultiValue technology into the mainstream computing market. With exclusive worldwide distribution rights, jBASE International offers technologies and assistance that allow businesses to thrive into the future.

Nebula R&D Announces NebulaPay

Nebula Research and Development has announced the release of the new NebulaPay financial payment processing interface for Pick/MultiValue business applications. NebulaPay provides a BASIC callable interface for processing credit and debit cards, address verification, check validation, gift cards, and Electronic Benefits Transfers. NebulaPay will support standard Point Of Sale (POS), Mail Order/Telephone Order (MOTO), restaurants with pre-auth and gratuities, Web shopping, and other consumer/vendor media.

NebulaPay is free from Nebula R&D, with no initial purchase or maintenance fees, and Nebula R&D will match or beat existing merchant fees. The architecture of NebulaPay allows for fast two-to-three second transactions with redundant servers over the Internet, with intelligent failover, load balancing, and other internal process management (developed by Nebula R&D) to ensure fast, stable, and secure transaction handling for small sites or large.

NebulaPay is now in beta over D3 and mvBASE, with jBASE and U2 soon to follow. Other platforms will be considered on request. Inquiries from VARs and end user/merchants are welcome: contact Tony Gravagno at (949) 380-1668, and visit http://Nebula-RnD.com/products/financial.htm

About Nebula Research and Development

Nebula Research and Development provides products and services which add value to MultiValue software by facilitating communications and interoperability with other systems and applications.

Microgen MultiValue Solutions Appoints NBA as mvQuery Reseller

NBA Consultants Ltd. (NBA), an independent company specializing in the provision of business solutions to the third-party warehousing and haulage industries, has entered into an agreement with Microgen MultiValue Solutions, a part of Microgen plc, to resell mvQuery, a MultiValue reporting solution.

NBA Consultants provides a comprehensive set of logistics management systems. Their approach is to provide every customer with uniquely tailored solutions based around a standard software package, thus providing full functionality with reduced maintenance overheads.

Microgen's mvQuery is a purpose-designed reporting tool for MultiValue environments, providing a complete solution for the design, generation and distribution of reports throughout an organization.

"After a protracted search over many years through differing products, we were extremely pleased to have found mvQuery," stated Tony Ledra, managing director for NBA. "MvQuery offers our customers a single reporting solution for ad-hoc reporting, embedded application form reports, data interchange and Web-based reporting."

"We are happy to welcome NBA to the mvQuery fold," said Andrew Muddiman, head of Microgen MultiValue Solutions. "NBA has proven over the past 11 years that a commitment to providing reliable applications of quality really does lead to a loyal customer base. We are proud NBA chose mvQuery as the only reporting solution that could complement their product set."

About Microgen

Founded in 1972 and listed on the London Stock Exchange since 1983, Microgen provides software, consultancy and IT services enabling businesses to collate, process and distribute corporate data to enhance their business processes and information output.

Over the last five years, the company has undergone a significant transformation from its original information management origins and is now a market-leading provider of IT software, solutions and services to major public and private sector organizations. **www.microgen.co.uk**

About NBA

NBA Consultants are specialists in storage and haulage solutions and have been operating in these fields for more than 11 years. Thanks to its consultancy and project management experience, NBA has extensive knowledge of operating procedures in these industry sectors and provides a number of specialist modules such as freight forwarding, excise and duty calculations, and interfaces to other systems. **www.nba-consultants.co.uk**

Revelation Software to Host User Conference

Revelation Software, a Westwood, N.J.-based software company specializing in application development tools, will be hosting a user conference June 24 - 26 in New Orleans, La. The conference schedule includes sessions for high-end technical attendees as well as sessions designed for the end users of Revelation-based applications.

Revelation staff will be providing live demos of their flagship product, OpenInsight. Revelation Partners along with industry experts will participate in the exhibitions.

The 2004 Revelation User Conference is an opportunity to learn the latest industry intelligence, find solutions to professional challenges, and network with the developer community. Almost 200 professionals from all over the globe will gather in New Orleans to attend education sessions, explore the latest technology in the field, and enjoy the best Mardi Gras Ball in years.

WHAT: 2004 Revelation User Conference

WHO: Revelation Software

WHEN: June 24 - 26

WHERE: The Crowne Plaza Hotel, New Orleans, La.

To register for complimentary admission to the exhibit hall, contact Nancy Ruane, nancy@revelation.com

U2logic Announces New Hires

U2logic, a leader in Web-enabling applications for the U2-driven business, has appointed Stuart Mackenzie to its Professional Services Group and Sharon Youngblood as its Key Accounts Manager.

Stuart joins U2Logic with excellent credentials, and eight years experience as a Senior Consultant with the IBM Professional Services Division. His expertise with SB+, MITS, U2 and RedBack will further U2logic's mission to provide high-quality, high-performance applications and easy-to-use tools that can help the U2-powered business realize faster Web-enablement, the company said.

Sharon joins U2Logic with eight years experience at UniData and Ardent as a Customer Service Representative and Sales Engineer. Her vast knowledge of infrastructure-based technologies and tools will allow U2logic to further assist MV-driven businesses pursue their goal of protecting legacy application investments.

With the increased demand for Web-enabled products and services, U2logic has expanded its new corporate headquarters, which is located at 8001 East 88th Ave., Suite 200, Henderson, CO 80640.

For more information on U2logic and how you can become Web-enabled and more on demand, visit its Web site at www.u2logic.com or call toll free at 1-866-XLr8me2 (866-957-8632).

newsmakers

Via Systems Signs Binary Star as South America Distributor

Via Systems Inc. has named Binary Star Development Corp. as the exclusive distributor of UniVision and ViaDuct fx to the Brazilian marketplace. UniVision is a significantly enhanced Pick R83, AP, and D3-compatible multiuser database management system, based on the MultiValue model, utilizing state-of-the-art object-oriented technology on Linux, Windows 2000/XP, SCO Unixware and AIX platforms.

ViaDuct fx is an industry-leading terminal emulation and MultiValue connectivity tool. Binary Star also announced the availability of its Nucleus AE product, a 4GL to UniVision.

"We are delighted with the response from the Brazilian marketplace to our product offerings," said Robert Catalano, president of Via Systems, "and from the opportunities that Binary Star is creating."

"UniVision coupled with Nucleus AE technology offers developers and VARs the ability to rapidly develop secure, easily maintainable GUI, character-based and Web applications, with a very low TCO (total cost of ownership)," said Lee Bacall, president of Binary Star Development. "Nucleus AE technology and UniVision offer simplified migration, a short learning curve and unparalleled flexibility.

"We have been gratified by the receptions we have received from each demonstration of Nucleus technology made to IT professionals in Brazil, and are excited about the potential for growth within the Brazilian marketplace," he continued. "We are looking forward to a long and mutually beneficial association between Binary Star and Via Systems.

Binary Star has recently appointed Porfirio (Matt) Sperandio, Pick veteran and long-time Binary Star staff member, as the director of Latin American Operations. Sperandio, an expert in system integration from Sao Paulo, Brazil, speaks many languages, including Brazilian Portuguese. He is working with developers, VARs and educational institutions in Brazil to demonstrate and promote the advantages of Nucleus AE Technology coupled with UniVision.

Sperandio shares a deep cultural affinity for the concerns of the Brazilian business and technical communities, as he is known among many IT executives and business associates alike. "We are happy to have Matt as a member of the Binary Star Development team as our director of Latin American Operations," Bacall commented.

www.binarystar.com

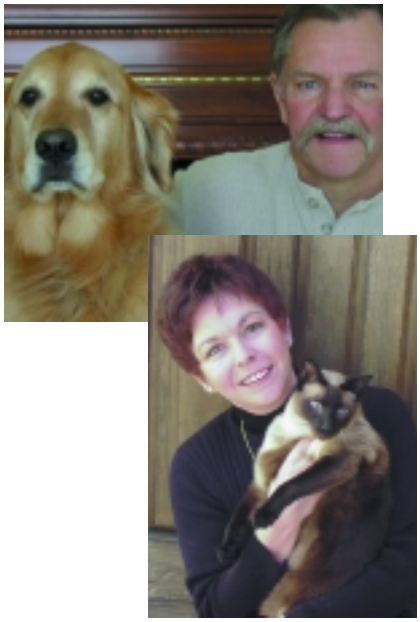

UniVerse and UniData Hashed FILES Part 2

BY JEFF FITZGERALD AND PEGGY LONG

Overflow: Causes and Performance Issues

In our previous article, *UniVerse and UniData Hashed Files — Part 1* (*International Spectrum*, March/April 2004, pp. 39-44), we introduced the basic concepts of hashed files, defined "modulo," "separation" and "block size" and demonstrated how a data record is assigned to a specific group based on the hashing algorithm and the record key. Further we described the condition that occurs when the data assigned to a group's primary buffer exceeds the buffer size. This condition is called overflow and has serious performance implications. Because overflow can have such a huge impact on system performance it is the sole topic for this part of our hashed file series. We will describe the three specific conditions that cause overflow, discuss the manner in which overflow arises from each condition, and explain how UniVerse and UniData deal with overflow internally.

A File Cabinet Analogy

In order to construct a conceptual picture, consider thinking of a database file as though it were a metal file cabinet. The folders represent the data records we wish to store. We can describe a 5 drawer cabinet as having a modulo of 5 and thus 5 groups are allocated to the file.

Think of each drawer as being a buffer. The separation or block size is the buffer size. It's the space in each drawer. There are letter size drawers and legal size drawers; at least in theory, the amount of "data" which can be stored in each is limited by the size of the drawer.

The type or hashing algorithm is that little label on the front of each drawer that defines the range of data that will be stored there, such as "A through G" in the top drawer. We know where to look for a folder that begins with the letter "B" because of the label on the front of the drawer. Our rule that only folders beginning with the letters "A" through "G" are to be stored in the top drawer means there is no need to search more than that one drawer. This is the real beauty and economy of hashed files — the location of a specified data record can be calculated. Only one group need be read to locate the record.

To recap our file cabinet analogy, the number of drawers is equal to the modulo. The size of each drawer is the block size or separation. The formula that dictates which "records" are stored in a specific drawer is the label on the front of the drawer — the hashing algorithm. This file cabinet analogy is simplistic but it works reasonably well to provide a mental picture of a database file. When a drawer in the file cabinet becomes completely full, you might place the excess file folders in a box to the side and put a note in the back of the drawer, saying, "Look in the box!" This situation could be termed "overflow."

Overflow is generally considered to be "bad." It can certainly increase the time required to access a specific data record. Imagine if you will that our "A" drawer is filled to capacity. Another large folder for the Aardvark Corporation needs to be added to the drawer. There is no space. Our database administrator (the filing clerk) adds a note to the drawer that she has stored this new folder in a cardboard storage box, 4th box in the 3rd stack in office 201. Suppose that you need this folder but don't know that it is in the overflow collection. You search the top drawer for Aardvark. It's not there. You discover the note that explains where you can find the folder. You proceed to that location and retrieve the folder. This is much the same process that is required when a database group has overflow. An offset defines the location where the overflow data are stored. An additional buffer has been added at the tail of the file and the offset defines the location. Had that Aardvark folder been in the first drawer you would have retrieved it much more easily and

Continues on page 36

UniVerse and UniData Hashed Files

Continued from page 35

quickly than having to search for it in the overflow storage area. This is the simplistic explanation of why overflow is generally considered "bad."

The Three Causes of Overflow

There are only three conditions that can cause overflow. The most common of these is that the total amount of data to be stored in a group is greater than the size of the group's primary buffer. For example, suppose a file is using a separation of 4 and a modulo of 1; this means the file has one group with a primary buffer that can hold 2,048 bytes. Further suppose that there are 12 records of 500 bytes each that must be stored in that one group, for a total of 6,000 bytes of data. If you have 6,000 bytes of data to store in the group, there will be about 4,000 bytes of overflow, because only about 2,000 bytes will fit in the primary buffer of the group! It will require two overflow buffers of 2,048 bytes each to store the additional 4,000 bytes of data. The group will then have one primary buffer and two overflow buffers. The three buffers are linked together via offsets. We describe this overflow condition as "the file is too small," meaning that the modulo, the number of groups, is simply too small to accommodate the amount of data to be stored.

The solution is an easy one. Just increase the number of groups by changing the modulo. Given our example of 6,000 bytes and a 2,048 byte buffer, it's obvious that we need at least 3 groups or a modulo of 3.

A second cause of overflow is that the distribution of data among the groups of the file is uneven. The file may have a large enough modulo and enough physical space but unless the data are distributed evenly, the space cannot be effectively used. One group may have too much data assigned and another may be very sparsely populated. We describe this condition as "a bad distribution." Sometimes we say the "type is wrong," meaning the choice of hashing algorithms could possibly improve the distribution (i.e., distribute the data more evenly among the groups).

In a paper filing system we can devise an almost infinite number of "formulas" or rules that assign the file folders to a specific location. For example, a drawer could be assigned labels "A-G," "AAA-ABA" or "A-CAA"; each label specifies a different range of folders to be stored in that drawer. Or, we might choose to file according to the zip code or customer number. We would hope for a scheme that would result in a "good fit" between the file folders and our filing cabinet. In our paper filing system, when one drawer begins to approach capacity we shuffle folders to the next drawer and re-label the drawers. This shuffling process is called "resizing" when applied to UniVerse and UniData files.

Unlike our paper filing systems, UniVerse and UniData have a limited number of hashing algorithms. We cannot create new formulas or rules to distribute specific data records to designated groups. UniData provides two hashing algorithms and UniVerse offers 17. All the hashing algorithms depend upon the record keys to determine where the data are stored — into which group each record will be written. Records keys (also called IDs) are often chosen because they have implicit meaning — customer numbers, zip codes, account numbers and social security numbers all have meaning. But they may also form unusual and unexpected distribution patterns; many more people have social security numbers issued in New York than do from Wyoming, for instance. When the hashing algorithm and the set of record keys interact to form an uneven distribution the probability of overflow is extremely high.

Large records are often overlooked as a source of overflow. Large records are defined as data records which exceed the size of the buffers used by the file — remember that the buffer size is defined by the separation and is also called the block size. We have often thought that a file statistic called Large Records which segregates the overflow caused by large records would be preferable to adding that species of overflow to the total reported overflow statistic. One of the many laudable design features of both UniData and UniVerse is the lack of a limit on data sizes; records, fields, values and subvalues can be of arbitrary length. These databases are excellent in their design to accommodate the variability in data record lengths. But, what is very attractive in the design also produces what we call "ugly" statistics. When large records are added to the total reported overflow, it is easy to view this statistic as an indication that a file is poorly maintained if one does not consider the reason for the overflow. If a file contains a number of large records it may appear to have an excessive amount of overflow. Large records are common in files containing a large number of multi-valued fields. They are also seen very frequently in cross reference data or indices — the cross reference record for "Smith" will likely be much larger than that for "Shostakovich"!

Another source of "large records" is seen in applications where notes, comments, and descriptions may be entered into a field in the database file. We are sure the designer believed these fields would be used for short notes of 50 characters or less. We once saw a case where a very conscientious data entry clerk had entered a 100-page legal document into a comments field!

Fortunately, both UniData and UniVerse are designed to allow for this sort of large record overflow without substantial degradation to the database performance. Unfortunately, both UniVerse and UniData do not report overflow due to large records as a separate type of overflow. Because this can create such ugly statistics, it often leads analysts to resort to extreme measures in an attempt to "get rid of this overflow" — the end result is often a file that performs very poorly.

Several years ago we were called by a UniData client because his application vendor was insisting that he get rid of ALL overflow in a particular file before the vendor would provide support for some issue which had arisen in the accounting system. Because he had been using our software to resize his files he turned to us for an answer. UniData allows a buffer size up to 16,384 bytes. On inspecting the file we found data records in excess of 50,000 bytes. The ONLY way to get rid of all overflow in this file was to delete all records greater than 16,384 bytes in length and increase the block size to 16. Additionally, if one of the two hash types provided by UniData couldn't distribute the data perfectly, the number of groups would have to be increased well beyond the amount of space needed to accommodate the total data size. Our very frustrated client didn't consider deleting data records as an option.

As you might surmise, arbitrarily increasing the block size (separation) and increasing the modulo to accommodate an inflated overflow statistic that is due to a few large records can have a detrimental effect on both performance and disk space utilization. Both UniVerse and UniData have clever designs that address potential performance issues caused by a small number of large records.

UniVerse and UniData Overflow Handling

Note that we are only describing "static" hashed files in this article. "Dynamically"

hashed files operate much like static files but the mechanics are a bit different. We will be devoting a future article to dynamic files and their parameters.

UniVerse database files use what we like to call the "traditional" method of linking overflow data to the appropriate group. When more data hashes to a UniVerse group than can be held in the group's primary buffer, another buffer of the same size is appended to the physical file. The overflow data are placed into the new buffer and the forward pointer of the last record in the primary buffer is set to reference the location of the data in the overflow buffer. If the overflow buffer becomes full, another overflow buffer is appended to the file and linked to the group. As more overflow buffers are used, the group becomes a chain of buffers. Reading this chain to locate a data record requires more system resources as the chain becomes longer.

UniVerse is unique in that it does not split a data record between two buffers unless the record qualifies as a "large" data record. Imagine a file using a buffer size of 512 bytes. This is called "separation 1" in UniVerse jargon. Further imagine that the average data record size is 300 bytes. Let's assume that the first data record written to a group of the file is exactly 300 bytes. Since the group's primary buffer is 512 bytes, this leaves less than 212 bytes free to hold more data. Suppose that the second record hashing to this group contains 250 bytes. It won't fit in the available space. We might expect that the last portion of the buffer would be filled to capacity by a fragment of the new record and THEN an overflow buffer would be attached to provide storage for the "tail" of the data record. NO. UniVerse will not split a data record across a buffer boundary unless it is larger than the buffer. UniVerse will pad the primary buffer and attach the overflow buffer. All of that second data record will be written into the overflow buffer. An offset connects the primary buffer to the overflow buffer.

This has some interesting implications. Generalizing from our scenario you could expect 40% to 50% of the disk space allocated to the file to be wasted space because of the average data record length. In this situation, we would recommend that the separation be increased from 1 to 4 so that each buffer could accommodate more than one average record. Not splitting records between buffers is good for data integrity. If a pointer is lost due to a system crash

Continues on page 38

Overflow Due to Large Records

How much of overflow is due to large records? If you want an easy way to check a data file, try this quick and easy technique. We are assuming that you can use the UniVerse line editor or UniData AE editor.

STEP 1 At TCL in the account where your "test" file is located, create an I-descriptor in the VOC file. Call it RANGE

Sample: **ED VOC RANGE**

```
0001: I
0002: INT((LEN(@RECORD)) / 1000)
0003:
0004:
0005: 10R
0006: S
```

This I-descriptor will get the length of each data record (LEN(@RECORD)), divide it by 1000 and drop the fractional piece of the answer. Implicit in the division is that we want to know how many records are in each range by thousands. For example how many records are between 1000 and 2000; 2000 and 3000, etc.

STEP 2 Create another record that we can use to count the number in each category produced by the RANGE I-descriptor. Call this record RANGE.COUNT.

Sample: **ED VOC RANGE.COUNT**

```
0001: I
0002: IF RANGE > 0 THEN 1 ELSE 0
0003:
0004:
0005: 4R
0006: S
```

This I-descriptor returns "1" if the range is greater than zero (meaning the record is 1000 bytes or larger. Zero is returned if the record is smaller than zero.

I-descriptors must be compiled before they are used. The CD, compile dictionaries, program expects the items to be in a dictionary file. Because we are going to store these items in the VOC we will need to trick the compiler. Create an entry in the VOC called CD.VOC.

```
0001: F
0002: VOC
0003: VOC
```

Because VOC is also in field 3, the compiler will assume this is the dictionary file and compile our I-descriptors. We used this command:
At tcl: CD VOC RANGE RANGE.COUNT
Are you wondering why we are adding these entries to the VOC and not to the dictionaries of the files we wish to inspect? Both UniVerse and UniData query languages search the dictionary of a file first for words in the query sentence. If the word is not defined in the file's dictionary, the query language assumes that it will be in the VOC. By adding these entries to the VOC we can use them to inspect any file in the account.

STEP 3 Create a query using the name of the file you want to investigate. Here is an example:
LIST myfile BY RANGE BREAK.ON RANGE TOTAL RANGE.COUNT DET.SUP

Your report will look something like this.

RANGE	RANGE.COUNT
0	0
1	175
2	74
3	49
4	20
5	5
6	9
7	4
8	3
9	3
10	2
15	1
	345

662 records listed.

Here's how we read this data:
Start with the second line that has 1 and 175 on it.

There are 175 data records that have between 1000 and 2000 bytes.

Then on to the next line:

There are 74 records with between 2000 and 3000 bytes.

There are 49 records with between 3000 and 4000 bytes.

There are 20 records with between 4000 and 5000 bytes.

And so on. In our example our buffer size, set by the file separation (block size), is 2048. Therefore any record greater than 2000 (rounding to make it simple) is considered a LARGE RECORD and won't fit in a single buffer. Therefore we recognize that there WILL be overflow in the file due to LARGE RECORDS.

Our conclusions are that there are (345 - 175) 170 records which fall into the LARGE category. (We subtracted the first 175 records because they are between 1000 and 2000, less than 2000 bytes.) The LARGE category is approximately 26% of the records in the file. (170/662) = .256. Obviously this will explain why our file has so much reported overflow. Is it poorly sized? Should we change the separation to "get rid" of overflow? This report does not address those two questions. The intent is to inspect the file to explain why there is a large amount of overflow.

There are three possible causes of overflow. These are 1) A modulo that is too small (i.e., not enough space was allocated for the file). 2) The TYPE is not distributing the data evenly over the available groups. Or 3) There are data records larger than the file buffer size (the separation or block size). Changing the modulo and type will not change the amount of overflow due to large records. Should you change the separation to "get rid" of overflow? In general we advise against it. In most cases increasing the separation higher than 4 will NOT improve performance and in many cases will be detrimental to performance.

UniVerse and UniData Hashed Files

Continued from page 37

during a file update, it will not cause part of a record to be "lost."

Continuing our example, let's consider the case where the second data record is 600 bytes in length. This record is larger than 512 bytes (the buffer size for the file) and is therefore considered a "large" record. UniVerse provides an elegant technique for storing "large" data records. In the case of a 600 byte record and a 512 byte buffer, the last 500 bytes of the data record would be stored in an "overflow" buffer. The first 100 bytes would be stored in the primary buffer (assuming sufficient space). A pointer referencing the offset to the overflow buffer is written in the record header along with a "large record" flag. The beauty of this design is that the overflow buffer

Looking for Overflow: A Few Useful Tools

You can determine whether a file has overflow by using the FILE.STAT command. The UniVerse version of this command produces a quasi-histogram on the last line.

GROUPS

25%	50%	75%	100%	125%	150%	175%	200% full
0	1	20	14	2	0	0	0

Any number below the 25% label can be read as X number of groups that are between zero and 25% full of data. The 100% column can be read as X number of groups between 76 and 100% full. Any numbers in the columns to right of the 100% column describe the number of groups that have overflow buffers. The highest percentage column is between 175 and 200% (meaning that there is a primary buffer and an overflow buffer — the chain is two buffers long). Every group with two or more overflow buffers will appear in the 200% column. The longer the overflow chain, in general, the slower the performance. From this utility there is no way to know the number of bytes in the largest group.

The UniData version of FILE.STAT reports the number of groups having Level 2 overflow.

File name = JTEST
Number of groups in file (modulo) = 1
Static hashing, hash type = 0
Block size = 1024
File has 8 groups in level two overflow. Please resize.

There is no histogram and no statistic reporting the largest group. The "guide" and GROUP.STAT utilities are also useful.

For UniVerse database files, the distribution of data across the groups can be inferred by mentally plotting the numbers on the quasi-histogram to form a curve. If the numbers on the ends (25% and 200%) are large it would indicate there are lots of groups that are sparsely populated and lots of groups that are in overflow. Calculating what we call the percent standard deviation of bytes per group is a more quantitative approach. The percent standard deviation can be seen in the statistics appearing at the end of the GROUP.STAT output, for both UniVerse and UniData.

It is often useful to determine the range of data records lengths in a data file. We have provided a few simple tools in the sidebar called "Overflow Due to Large Records." FILE.STAT and HASH.HELP.DETAIL (UniVerse only) report the number of bytes in the largest record and the average number of bytes. This is insufficient information to determine if most of the overflow in the file is due to large records.

Have you concluded that studying file statistics can be very time consuming? We agree with you. However, to devise a resizing strategy to improve performance and minimize overflow you need to know the cause of the overflow.

is completely filled, with no wasted space. Additionally, the only time the overflow buffer will be read is if/when the large record is read. The large record data is not read when scanning the group and doesn't impact performance as much as overflow due to non-large records.

UniData has a unique design that separates the record keys and record data for all records. This cleverly accommodates large records as well as a modest amount of data overflow without significant performance degradation. Visualize a buffer with a point marked near the center; we will label this point "DATA.POS." Now think of the left side of this point. This is the relative position where the first record key written into the group will be stored. To the right of DATA.POS the data will be stored. An offset that points to where the data begins is written at the top of the buffer. As data are added to the file, the offsets accumulate to the right of the previous offset and the keys are added to the left of the last key written into the buffer. Data records accumulate to the right of the previous record.

The primary buffer will continue to accumulate data until the space reserved for offsets and keys is filled. At that time UniData moves the data in the primary buffer to overflow buffers and utilizes the vacated space in the primary buffer for more keys and offsets. The overflow buffers occupied by data are called "Level 1" overflow. In general, the data portion of the buffer fills before the top portion containing the offsets and keys.

As more records are added, the primary buffer can become full of record keys and offsets. When this occurs UniData appends another overflow buffer and treats it as though it were a primary buffer, placing new keys, offsets and data into it. The primary buffer contains a pointer to this new overflow buffer containing keys, offsets and data. This is called "Level 2" overflow. (If you are interested in a more extensive discussion of Level 1 and Level 2 overflow, send an email to peggy@fitzlong.com or jeff@fitzlong.com. We will send you an Acrobat file containing a paper we wrote called "UniData Files: Level 1 and Level 2 Overflow Explained". If you have a copy of our FAST User's Guide, this paper is in the appendix, p.16.)

This is somewhat similar to UniVerse's handling of large records; the overflow data doesn't have to be read while scanning the group for a specific record. Each data record is located, when required, via the offset in the top of the primary buffer. However, Level 2 overflow is a linked list and the chain of buffers containing keys and offsets must be read sequentially when scanning the group, with a substantial impact on performance as the number of buffers in the Level 2 overflow list grows. **is**

Overflow due to insufficient space — our "too small" category of overflow is simple and easy to fix. Resize the file with a modulo that produces enough groups to store the data. We suggest adding approximately 20% extra space for growth.

Overflow due to poor data distribution can sometimes be solved by changing the hashing algorithm to one which more evenly distributes the data among the groups of the file. UniData has only two choices and, unfortunately, no utilities to test which of those choices would be better. You must create files and compare the statistics. This can be very time consuming. UniVerse has the HASH.AID utility, which allows you to simulate resizing the file using all 17 types. (Use the syntax "HASH.AID filename 2,18,1 * * NO.PAGE" and then "LIST HASH.AID.FILE" to see the results. Clear the HASH.AID.FILE between files or the results will be confusing.) This is a good way to choose the best type parameter, but it too is very time consuming. Recognizing the best selection can be tricky because the utility does not report the number of empty groups or the percent standard deviation. We don't like using HASH.HELP (UniVerse only) to determine file types — this utility tries to categorize the record key structure and is often wrong — a simulation method is much better.

Overflow due to large records may not need to be "fixed." Both UniData and UniVerse store large records cleverly in order to minimize degradation in performance due to a few large records in a file. We suggest using our simple I-descriptors described in the sidebar to evaluate how much of the overflow is due to large records. You may want to view the data in the largest of these to determine if the data are worth storing. For example, if the largest record in a cross reference file is indexing all of the occurrences of the word "the" in company names, your best solution is to determine if there is a way to eliminate indexing words like "the," "and," "Inc.," "Co," and "A."

Overflow can cause poor response time in data retrieval. The degradation in performance is due to the increased system resources required to read and write overflow buffers that are linked to the primary buffer. As the chain of overflow increases in length, the performance impact increases geometrically. Minimizing overflow can improve performance by reducing the number of buffers that must be read to locate the desired data or that must be written during an update. It is well worth the time to systematically analyze and resize UniVerse and UniData files.

Knowing the file structure and how overflow occurs will guide us when designing files — we will create files that behave well. However, many of us have purchased or inherited files that someone else designed and which may not be perfect in terms of controlling the factors that cause overflow. These files may not be well distributed by any of the available hashing algorithms. And they may contain large records and widely varying record sizes. Attempting to remove all overflow from such files is often an exercise in futility; extreme measures, such as drastic modulo and/or block size increases, will often result in worse performance than if we ignored the file!

In summary, overflow can cause performance degradation. And, by resizing the file to reduce overflow caused by a modulo that is too small or a poorly chosen hashing algorithm, you can expect to improve system performance. Substantially increasing the modulo or block size to reduce overflow due to large records or lack of an effective hashing algorithm is unlikely to improve performance and may possibly cause the opposite.

Twenty years ago Peggy Long and Jeff Fitzgerald were running a critical benchmark on a top-of-the-line Prime INFORMATION system. The benchmark aborted. After several hours of detective work they identified a damaged file. That started a discussion concerning performance, broken files and how to check the internal structure of files after a system crash.

After several weeks of work using Peggy's FORTRAN skills and Jeff's INFO BASIC knowledge they were confident that they understood the file internals. This led to a utility that would quickly scan a file, report errors and recommend the optimum MODULO and TYPE parameters needed to RESIZE the file. A year later they began marketing FAST, which evolved from this utility.

BY RAY ELSE, I.B.M.

PHP Web Development

Using RedBack Objects

Hypertext Pre-Processor (PHP) is a popular new Web middleware development language that you might want to use one day to access your IBM UniData or UniVerse database. Most popular on Linux, it also runs on Unix and Microsoft Windows Web Servers. This article shows how you can use PHP to access your UniData or UniVerse database in a simple, elegant way using IBM's RedBack middleware.

Why use RedBack and RedBack Objects with PHP? Well, besides RedBack being a solid performer now for over five years, with thousands of installs all over the world, it also allows you to access UniData or UniVerse without having to write in any other language except PHP in the middle tier and Basic on the database side. And RedBack Objects are virtual/reusable — if tomorrow you decide to use Microsoft Active Server Pages (ASP), ASP.NET, Java Server Pages (JSP), or Cold Fusion instead of, or in addition to, PHP, you won't have to change your RedBack Objects at all!

Similar to ASP or JSP, PHP is a middleware engine that works with your Web Server to dish out dynamic Web pages. A PHP programmer adds PHP script to Web pages that the PHP engine will interpret. This script can do powerful things, like make calls to databases to allow the insertion of database data into a page, or to write data away.

The way a Web Server like Apache or Microsoft IIS handles requests for PHP pages is that if a requested page's name ends in ".php", then the Web Server calls the PHP program before it sends the page to the requesting client browser. The PHP program processes the embedded PHP script in the page and returns the HTML presentation results as a page to the Web Server to be sent to the requesting client browser. (Note: the Web Server needs to be configured to recognize the php suffix as a trigger to call the PHP program.)

PHP has extensions allowing calls to other language objects. For example, there is a Java extension that allows a PHP programmer to instantiate a RedBack RedObject to set properties, get properties and call methods on one of your RedBack Objects. The PHP programmer doesn't need to know Java at all to use the RedBack RedObject class to use RedBack Objects.

Below is a very simple PHP page called test.php. If requested from an IIS or Apache Web Server that allows scripts to run, this page should print "This is a PHP Page" and php environment details.

```
<html>
<head><title>Test PHP</title></head>
<body>
This is a <?php print("PHP") ?> Page
<p>
<?php
phpinfo();
?>
</body>
</html>
```

The Url input on a browser to request this page would look something like this:

```
http://web_server_name_or_I_P_address/directory/test.php
```

Note the <? ... ?> script tags, similar to ASP and JSP's <% ... %> tags. The code in these tags are processed and the resultant HTML page sent to the browser with the script removed.

PHP IMPLICIT OBJECTS

PHP comes with implicit objects similar to ASP and JSP for facilitating access to submitted form data, cookies, etc.

_COOKIES - array of cookies; the keys in the array are the names of the cookies

_REQUEST - the requesting page; form and Url and cookie name/value pairs. Combines the contents of other superglobals like _GET, _POST, _COOKIES, and _FILES (warning - where variables have the same name they can be overwritten)

_SERVER - array contains info describing the server and its environment, with elements like DOCUMENT_ROOT, QUERY_STRING, REQUEST_URI, HTTP_REFERER

_SESSION - array contains variables placed in a user's PHP session; like ASP and JSP. This session is based on a unique Id stored in a memory cookie of the client browser; note you may have to change the session.save_path in php.ini, and you may want to set session to be on by default by setting the session.auto_start in php.ini.

PHP SYNTAX

The PHP scripting language is similar to perl (which is similar to Javascript). Although this guide is not intended to cover the rich syntax of PHP and all its objects and extensions, the following are handy reminders of the syntax. Single line comments are done with two slashes or the pound sign.

```
// this is a comment
```

Concatenation is done with the dot character (no spaces required).

```
$myvar = "the " . "end";
```

Variables begin with a $ and are not typically declared. PHP will type them as required. The scope of a variable is confined to the code block in which it was declared ("if" statement, loop, function). The following are example assignment statements:

```
$mytext = "123";
$myinteger = 123;
```

The operators are the usual: **+, -, *, /, %, +=, -=, ++, --**

PHP supports single and multi-dimensioned arrays. Arrays in PHP can be indexed by integers or can be associative arrays (hashes). A single dimensioned associative array can be thought of as a row with named fields. You can initialize an array by setting each element:

```
$name[0] = "Ray";
$name[1] = "Archie";
```

or

```
$sport["ray"] = "football";
$sport["bob"] = "hockey";
```

We can then reference an element in the array so:

```
$strName1 = $name[0];
$strName2 = $name[1];
```

or

```
$strSport1 = $sport["ray"];
$strSport2 = $sport["bob"];
```

Two functions of interest with arrays are explode and implode. The explode function turns a delimited string into an array while the implode turns an array into a string.

Example using explode:

```
$VM = chr(253);
$myVMstring = "Ray".$VM."Archie";
$myarray = explode($VM,$myVMstring);
print('$myarray[1]='.$myarray[1]);
```

The conditional operators are like Javascript: **==** (two equal signs for comparison), **!=** (not equal), **>, <, >=, <=, !, &&, ||, xor.**

```
if ( $init == "RE" || $init == "REE") {
   $fullname = "Ray Erstwhile";
} else {
   $fullname = "John Jenkins";
}
```

For looping, one can use a While loop, Do loop, For loop, or Iteration Loop:

```
for ($cnt = 1; $cnt <= 10; $cnt++) { print("$number<br>\n"); }
```

An iteration loop looks so:

```
$sport["ray"] = "football";
$sport["bob"] = "hockey";
foreach($sport as $key=>$value) {
print("sport " . $key . " is " . $value . "<br>
");
}
```

PHP has hundreds of built-in functions. You can also write your own. Parentheses are required when calling functions.

REDBACK OBJECTS

With the RedBack Object Designer tool, one designs RedBack Objects (RBOs) with properties and methods. The properties are typically data place holders (like variables) whereas the methods are typically pointers to IBM UniBasic subroutines. The principal RedBack objects one can use as templates for creating one's own objects are the following:

- **SLRBO -** the **fastest** RedBack Object, this very flexible, stateless object is good for single record reads (or multiple records delimited in some way) and writing away data to one or more files; requires you to write the methods/subroutines that do the reading or writing. We recommend you use this object whenever possible in your application as it has the lightest load on a system.

- **RBO -** this flexible but less fast **stateful** object is good for the same tasks as the SLRBO, but being stateful it can also be used to keep data around between page requests. Typically one stores off the handle to the object in a cookie or a session variable so one can refer to it again later by opening it with this handle. Note: an alternative to using a stateful RBO to keep data around would be to use an SLRBO along with Session Variables.

- **uObject -** will do simple reads and writes for you without you having to write Basic code; has optimistic locking support; less flexible and less fast than RBO or SLRBO though, and more complex to work with

- **uQuery -** stateful report object; given a U2 (pick flavor) Select Statement it returns a recordset. You don't typically write Basic code with this object. One can keep a handle to the object and page through the data on multiple hits.

REDBACK SETUP

RedBack setup on the Database Server and on the Web Server is covered in the RedBack documentation. As part of that setup, a test database called rbexamples is installed, complete with data and RedBack Objects for accessing that data. We will access rbexamples using those objects in our example code to come.

Since we are going to be making calls from PHP to RedBack Objects from Linux and/or Windows, we will use the RedBack Java component RedBeans.jar on the Web Server along with the rgw program and the rgw ini file to access the rbexamples database:

1. **RedBeans.jar** - contains the RedObject class that enables us to use our RedBack Objects; place anywhere you want on Web Server, we will then make an entry of it in our php.ini file (see next section PHP SETUP WITH JAVA EXTENSION).

2. **rgw.dll** (on Windows, rgw on Linux) - we will place this ISAPI program in the scripts directory of our IIS Server on Windows (on Linux/Apache we would place it in the cgi-bin directory), and use it to connect from our RedBeans' RedObject class to our database to use our RedBack Objects (examples to follow).

3. **rgw.ini** - we place this ini file in the root document directory. On Windows this is typically wwwroot, on Linux/Apache it could be htdocs; this file contains the database name followed by a space followed by the database server machine name or i.p. address, a colon, and the port number where the RedBack Responder is listening for requests to this database (i.e. - rbexamples my2000:8444).

PHP SETUP WITH JAVA EXTENSION

PHP setup on Linux or Windows is covered at www.php.net and several other sites you can search online. We will assume here you have PHP installed and now you want to get the Java extension working so you can make calls through RedBack.

Before you try to use Java classes (like those in RedBack's RedBeans.jar file), you will need to download a jdk (containing a jvm [java virtual machine]) onto the Web Server (if you haven't already got it).

continues on page 42

Continued from page 41

On Linux you compile PHP with Java support (again search online for examples). On Windows you simply edit the php.ini file.

NOTE - the following instructions are for Windows, but they would be similar for Linux.

Step 1: After installing the JDK, add its bin directory to the environment PATH. Note that JAVA_HOME and CLASSPATH system environment variables are not used by PHP but instead you must edit these variables in the php.ini file - see step 2.

Step 2: Modify your php.ini for Java and while you are at it, add the RedBeans.jar to the path. (On Windows you would have copied the php.ini-recommended file to Windows system root, for example WINNT on Windows 2000, and changed its name to php.ini.)

a) uncomment the line

```
extension=php_java.dll
```

b) Under the [Java] declaration, set up your paths, for example this is what they look like on my PC running Windows 2000 with the RedBeans.jar in my RedBack directory on the C drive:

```
[Java]

;R.E.,2/17/2004 set php java extension jar and
RedBeans.jar in java classpath

java.class.path = "C:/php4/php-4.3.4-
Win32/extensions/php_java.jar;C:/RedBack/Red-
Beans.jar"

java.library.path = C:\php4\php-4.3.4-
Win32\extensions\

java.home = c:\jdk1.3.1_06

java.library =
c:\jdk1.3.1_06\jre\bin\hotspot\jvm.dll
```

Note the Java.class.path should have both the php_java.jar path and the RedBeans.jar path.

Step 3: Test your Install by calling a simple Java class. Create the following php page and test it from the browser:

```
<?php

$system = new Java("java.lang.System");

print "Java version=".$system-
>getProperty("java.version")."<br>\n";

print "Java vendor=".$system->getProperty("java.ven-
dor")." <p>\n\n";

print "OS=".$system->getProperty("os.name")." ".

$system->getProperty("os.version")." on ".

$system->getProperty("os.arch")." <br>\n";

$formatter = new Java("java.text.SimpleDateFor-
mat","EEEE,

            MMMM dd, yyyy 'at' h:mm:ss a zzzz");

print $formatter->format(new Java("java.util.Date"))."\n";

?>
```

You should see something like:

```
Java version=1.2.2
Java vendor=Sun Microsystems Inc.
OS=Windows 2000 on x86
Friday, Feb 20, 2004 at 10:22:45 AM Standard Time
```

Assuming that your Web Server can "see" your Database Server (and if a firewall exists between the two, you have opened up the port [as shown in rgw.ini on Web Server and rgwresp.ini on Database Server] for rbexamples database), you are now ready to access the UniData or UniVerse rbexamples database.

DATA DISPLAY EXAMPLE

The following PHP page prompts for an employee Id (1001 - 1025 are valid), makes a database connection, creates an object, sets a key property, calls the database read method of the RedBack object, then displays the returned properties' values, the first name and last name of the employee from the EMPLOYEES file in the rbexamples database.

NOTE: this stateless example will ONLY work with an SLRBO. To make it work with an RBO or a uObject you would need to issue an "open" on the object — see the comments in the code. A stateful uQuery (which returns a recordset) is different still — please see the REPORT EXAMPLE.

```
<?php
// stateless.php, Ray Else, IBM, 2/20/2004
// RedBack stateless object example - no warranty implied
$empId = "";
$firstname = "";
$lastname = "";
if (isset($_REQUEST["Id"])) {
 $empId = $_REQUEST["Id"];
 if ($empId != "") {
  // RedBeans.jar needs to be in php classpath in php.ini
  $rbObj = new Java("com.ibm.redback.redbeans.RedObject");
  // below we are using IIS ISAPI program to connect, for Apache
  // we would use a string like "http://localhost/cgi-bin/rgw/rbexamples"
  $rbObj->setRBOAccount("http://localhost/scripts/rgw.dll/rbexamples");
  $rbObj->setRBOClass("EXMOD:EmpReader");
  // if using stateful object like uObject or RBO would need to open or refresh it like so:
  // $rbObj->open();
  // but since we are using stateless object here we can immediately use without open
  $rbObj->setProperty("EmpId",$empId);
  $rbObj->callMethod("DoRead");
  $firstname=$rbObj->getProperty("FirstName");
  $lastname=$rbObj->getProperty("LastName");
  $rbObj = NULL;
 }
}
?>
<html>
<head><title> Display Employee </title></head>
<BODY onLoad="javascript:Form1id.focus()">
<h2> Display Employee </h2>
<form name="Form1" method="POST"action="stateless.php">
Enter Employee ID: <input type="text" name="Id" value="<?php print( $empId) ?>">
<input type="submit" value="Read"><p>
First Name is <?php print( $firstname) ?><br>
Last Name is <?php print( $lastname) ?>
</body>
</html>
```

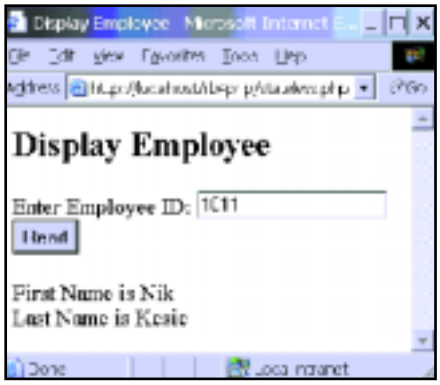

REPORT EXAMPLE

If you can gather the records required for a report using a single select statement, then the uQuery object is an ideal RedBack Object to use to bring the data from the database. This object returns a record-set, which we can loop through to display the fields in each record.

```
<?php
// report1.php, Ray Else, IBM, 2/17/2004
// RedBack uQuery object example - no warranty implied
$pageCount = -1;
$recCount = 0;
// RedBeans.jar needs to be in php classpath in php.ini
 $rbObj = new Java("com.ibm.redback.redbeans.RedObject");
// below we are using IIS ISAPI program to connect, for Apache
// we would use a string like "http://localhost/cgi-bin/rgw/rbexamples"
$rbObj>setRBOAccount("http://localhost/scripts/rgw.dll/rbexamples");
// want to use uQuery object EmployeeList in EXMOD module
 $rbObj->setRBOClass("EXMOD:EmployeeList");
  // if using stateful object would need to open or refresh it like so:
  $rbObj->open();

// could set select_criteria property but in this case we want all records
// $selectstmt = some valid select statement
// $rbObj->setProperty("select_criteria",$selectstmt);

$rs = $rbObj->callMethod("Select");
$pageCount = $rs->getMaxPages();
$recCount = $rbObj->getProperty("MaxRows");
?>
<html>
<head>
<title> Employee Report </title>
</head>
<BODY>
<P>Records Found = <?php print( $recCount) ?>
<table border=1>
<tr>
   <td><b>No.</td>
   <td><b>Last Name</b></td>
   <td><b>First Name</b></td>
</tr>
<?php
if ($pageCount > 0) {
  $recNumber = 0;
  $rs->moveFirst();
  for ($ii = 1; $ii <= $recCount; $ii++)
  {
   $recNumber++
?>
     <tr>
    <td><?php print($recNumber) ?></td>
          <td><?php print($rs>getProperty("LAST. NAME")) ?></td>
          <td><?php print($rs->getProperty("FIRST. NAME")) ?></td>
     </tr>
  <?php
  $rs->moveNext();
   }
  }
  ?>
  </table>

  </body>
  </html>
  <?php
  $rbObj = NULL;
  $rs = NULL;
  ?>
```

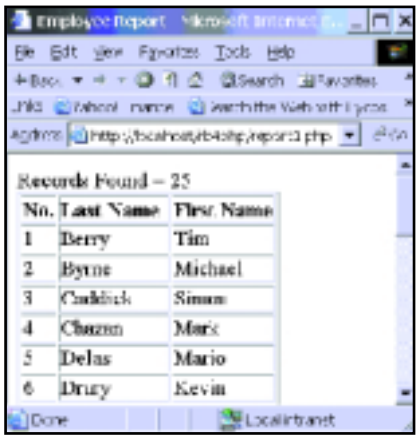

SUMMARY

In this article we have covered how one could write a PHP Web Application using RedBack to access a UniVerse or UniData database. If you have any questions, please feel free to contact raymonde@us.ibm.com or IBM U2 sales or IBM U2 support.

It's Time to Graduate

The Internet Needs to Grow Up a Little

BY MELVIN SORIANO

THE OTHER DAY I GOT A GRADUATION announcement from Trojan Giobbi, a proud moment indeed. As we watch kids grow up into adults, we also watch as they move on (and likewise, stop sending bills home). Well, it sure seems like the Internet needs to move on a little. We've all certainly watched it grow up: learning to talk with normal people, wild and crazy teen years, slow integration into the business community. But sometimes you just can't wait to see it grab that diploma and act like a grown up.

Take email, for example. Love the email. People can update you, cajole you, quiz you, razz you. It's globally used and appreciated. You can access it anywhere it seems, and not just on your desktop. Cellphones have email. Starbucks has email. Hotels have email. If my gym put in Web browsers on the treadmills, I'd look like a California governor in no time.

But it's getting tougher to reap the total rewards of email because with it comes so many other problems and frustrations. Okay, so you don't have to buy it health insurance, but email carries its own ills. Bugs and viruses, worms, and spam continue to saturate our emails and invade our computers. These abominations have moved from nuisance to expensive leaches on the corpus technologica. As they invade sophisticated network systems, they're turning a delicate yet productive doily that is the Internet into a crazed slab of silly string.

The rise of bad boys on the Web has been increasing for some time, true, and I'm certainly not the first person to threaten to use his PC as a binary paperweight. Still, the ever increasing assaults on networks and network security should be a wakeup call for even the Internet's most ardent fans. The powers that be that help organize the Internet are not only amenable to more heavy-handed intervention, they're actually starting (gasp) to email each other about it.

Some ideas that have been put forward are surprisingly aggressive. For example, some people are pushing for laws that permit customers to sue software companies over loopholes in their products. This is a big change; the software industry is currently exempt from such claims for the most part. Some are proposing to build a tracking system that supposedly would make it impossible for anyone, including hackers, to use the Internet without leaving an audit trail. A milder idea (probably supported by Symantec and McAffee) would be to force all PC users to install security.

Every idea has its problems, of course. Most obviously, they increase costs or reduce privacy. More subtle a change, however, would be the change to the Internet; they will be difficult to institute without changing the way the Internet, and the people who make the Net happen, work to date. It does seem clear, though, that the Internet has grown way beyond its nerdy origins. It's not a utopia after all, and if so, it needs to be managed to keep the bad from overtaking the good.

And the bad is really bad. Last year's attacks weren't just limited to desktops but were also pointed at central Internet DNS servers. Yeah. Nasty. Surprisingly, the worms haven't brought the entire Net down, which I suppose endorses the strength of the distributed design. But it sure leaves us feeling like Net performance can be flaky at best, and vulnerable to malice always.

Of the various proposals bandied about, the most radical technical idea would change the way data is transmitted on the Internet. The packets would no longer, the idea goes, be totally anonymous. Let's face it: the secrecy creates crevices for hackers so that they can forge or spoof the IP address on their data packets.

Bulk emailers (okay, spammers) exploit the secrecy obviously, as it masks them as they send zillions of emails with fake addresses. Many propose building mechanisms for the identification and authentication of everyone who uses the Internet. It basically exposes the dirty old man in a trench coat trying to flash you his email.

Probably a combo of two new standards will do much to limit these anonymous packets. Ipv6, the latest Internet protocol, allows assigning tags to data that move over the Net. If combined with DNSec, a more secure form of DNS, you'll be doing much to limit hidden attacks. DNSec uses server-based digital certificates to validate that the data passing across the wires originated from servers that they claim to be coming from.

Other technologies can also support these goals, but they'd require far more changes to the Internet's infrastructure. For example, requiring that every device connected to the Internet requires a certificate — that's a lot of management, which means costs, which means probably a no go.

Trial lawyer-types push for changing the laws. Making companies liable for their software would vastly improve software according to them. This is the consumer liability angle. It would make the Internet safer, and give all software companies more incentive to release bug-free software. Banks and financial institutions would also be made liable for the release of private information on the Internet.

If this happens, expect software upgrades to cease as we know it. Very few software companies will release any versions of their software without testing it on rabbit eyes and baby seals. Oops, they'll fall afoul there too. The point is, the tremendous pace of software development would be totally impaired if more time were spent worrying

about lawsuits than on product features. And comparing an operating system to a toaster might sound legally just, but the toaster does not have to be backwards compatible for 20 years of high-speed innovation in software and hardware.

And it's not just software programmers that would have to suddenly achieve discipline. It would be the users themselves. Even in today's lax security and software implementation world, users persistently complain about the draconian system admin who won't let them install their new software just yet. Will these complainers, that is, you and I, be upset if the delays were even more pronounced? Or do companies really want to wait three months for a product release instead of one week?

On top of all that, new laws and regulations might work locally, but the Internet is an international system. Local companies would be hurt as foreign companies would be able to evade prosecution and lawsuits for far longer.

So, the reality is that adding new restrictions is always a double-edged sword. Look at even the simple practice where ISPs block email that wasn't sent directly though the ISP's own email server. Spammers do this to blast millions of messages without being caught. But many people use dial-up to connect to corporate email servers or to connect to more secure servers than even the ISP might offer. Or some ISPs accidentally block bulk email from people who are trying to conduct legitimate communication with their customers.

Of course, the biggest problem won't be the businesses but home consumers. Most of the viruses that attacked last year were launched by home computers connected by high-speed connections. In other words, the people who least likely understand security or care about security of other computers are quite responsible for the problems. Of course, consumers care about their own computers, but they will certainly question why they should worry or pay to help keep AOL and Microsoft secure.

Some say that ISPs should require consumers to install and subscribe to security software, and to purchase firewalls and secure hardware. Right. Unless that firewall can play MPEGs and play shoot-em-up, you won't see too many home buyers plunking down hard cash willingly. But some ISPs, AOL included, are starting to offer additional security, at a cost. But security hawks won't be satisfied until every PC is inoculated, just like children are expected to be before they can attend public schools.

And all this is to stop the simply vandalous. What about those with evil in mind? If all these hacking and cracking energies were turned into devious efforts, then personal identity security and national security would most certainly be compromised.

In the end, it's all a matter of maturity. We're all expected to play well with each other, grow up, find out what we're all about, and do what we do best. Massively integrative and communicative media like television and radio have matured; the Internet needs to do so as well. And we should insist that it happen and never let it falter. We as Internet users need to start setting our expectations on what a mature world would be like, push the Internet in that direction, and at all times, fight on. **is**

Melvin M. Soriano works at Eagle Rock Information Systems (ERIS), an Internet Application Service Provider and WebWizard/MultiValue Developer. ERIS has deployed enterprise-wide solutions on most MultiValue platforms and operating systems. HTM-Mel can be contacted at mel@eriscorp.com and visited at http://www.eriscorp.com/. You can always call him directly at ERIS's Pasadena offices: (626) 535-9658.

continued from page 13

Attribute 1 above contained "ABC]DEF]GHI]JKL]MNO", then the result displayed would have 5 rows. (i.e., Number of Items multiplied by Maximum Number of Multivalues).

ID	ATTR1	ATTR2	ATTR3
ITM01	ABC	A002	123
ITM01	DEF	A002	456
ITM01	GHI	A002	789
ITM01	JKL	A002	
ITM01	MNO	A002	

5 rows listed.

If there are no corresponding multivalues in other multivalued attributes, then NULL values are displayed for the corresponding multivalued positions as in column 'ATTR3' above.

The table is always exploded even when selecting non-exploded columns.

```
SELECT ID, ATTR2 FROM T1
```

ID	ATTR1
ITM01	A002
ITM01	A002
ITM01	A002

How you restructure the SQL catalog to handle a multivalued attribute depends on the type of data in the attribute. Examples of different types of multivalued data are listed below in the PERSONNEL file item.

001	*Person-id*
001 MR	Title
002 MARK]ROBERT	*Forenames*
003 ROBINSON	*Surname*
004 1 HIGH ST] STEVENAGE]SG4 1RE	*Address*
005 MCSE]INTRO TO IMPROMPTU	*Qualification*
006 11110]11182	*Qualification date*
007 MICROSOFT]COGNOS	*Qual. Company*
008 10959]9951	*Post date*
009 AC1]RE2	*Post company code*
010 10166]10023	*Post Department code*
011 P12]P54	*Post pay code*

Note: "]" denotes a value mark.

Techniques for Handling Multivalues

The techniques used to handle these multivalued attributes depend on the type of multivalued data. See below for a list of possible techniques and associated examples.

Attributes 2 and 4 to 11 all contain multivalued data.

Attributes 2 and 4 contain multivalued data, which needs to be on a single row.

Attributes 5 to 7 and 8 to 11 comprise sets of multivalued attributes containing related data. These attributes could be treated in a similar way to the single-row multivalues. However, a better technique is to create a "virtual" secondary table containing only the related data.

Technique Operation

1. Set up a column to extract the first value and ignore the rest.
2. Set up a column to extract the n'th value and ignore the rest.
3. Create a single field by concatenating all the multivalues together, separated by spaces.
4. Create a single field by concatenating all the multivalues together, separated by commas.
5. Expand the multivalues, one per row.
6. Multiple rows of multivalues containing data which is linked together can be handled by creating a 'virtual' secondary table in the main table to contain the data with one row per multivalue.

Examples of Handling Multivalued Data in a Single Row

Attributes 2 and 4 in the example item shown contain multivalues that need to be displayed all in one row.

Example 1—Attribute 2 Forenames

Attribute 2 in the PERSONNEL item contains the forenames of a person, separated by value marks. For example:

1. If only the first name is needed, set up a column to extract the first forename and ignore the rest.
2. If all the person's forenames are needed, concatenate the names together, separated by spaces.

Example 2—Attribute 4 Address

Attribute 4 in the PERSONNEL item contains the address of the person, with each part of the address separated by value marks. For example:

Set up a column for each address value, e.g. Address Line 1, Address Line 2, etc., using Technique 2 to extract the required multivalue in each column definition.

Examples of Handling Related Multivalued Fields

Attributes 5 to 7 and 8 to 11 are two sets of multivalued attributes, which contain related data. The first set (5 to 7) contains data which describes the qualifications held by the person, and when and where they where obtained. The second set (8 to 11) contains details of posts held by the person both, previous and current.

The recommended method of dealing with multivalues of this type of data is to create a "virtual" secondary table derived from the main table, (Technique 6). Exploding the multi- or sub-valued attributes of the main table will populate this table. The combination of the item-id and the multivalue (or subvalue) number can be used as the unique primary key of such a table.

We now have "normalized" data, which can be accessed and updated from external sources using ODBC and/or JDBC, using thousands of software products such as data mining tools, etc.

Summary

So in Reality, "two way" relational access, natively embedded within the product to any ODBC-compliant database, is possible and available now. For more details, see http://www.northgate-is.com/reality. **is**

northgate

101 Academy

Suite 200

Irvine, CA 92612

Tel: 866 473-2588

Fax: 866 823-1010

Main 866 823-1000

Email: realityusa@northgate-is.com