

INSIDE! BUSINESS KNOWLEDGE FOR PROGRAMMERS CONTINUES

\$7.00 U.S.

I N T E R N A T I O N A L
SPECTRUM®

THE BUSINESS COMPUTER MAGAZINE

MAY/JUNE 2006 • AN INTERNATIONAL SPECTRUM, INC.



PUBLICATION

Digging
through the
DATA
OLAP and Pivot Tables

PLUS! Getting
ready for SOA



Innovations by InterSystems



Rapid development with robust objects



Lightning speed with a multidimensional engine



Easy database administration



Massive scalability on minimal hardware

Database With Multidimensional Appeal.

Only Caché offers native object, robust SQL, and direct access to a powerful multi-dimensional database. And now, Caché has been enhanced with MultiValue technology, providing built-in support for MultiValue Basic, data structures, and queries. This enables MultiValue developers to preserve their existing application investments, plus make use of the latest Web, object, and SQL technologies.

By delivering new levels of high performance and rapid Web application development, Caché has become the world's fastest growing database product. Caché easily handles the most demanding applications with scalability from one to more than 50,000 users, with modest demands for hardware and maintenance.


We are InterSystems – an innovator in database systems for over 25 years, providing 24x7 support to 4 million users in 88 countries.



www.InterSystems.com/mvWW

I N T E R N A T I O N A L
SPECTRUM®

M A Y / J U N E 2 0 0 6



Digging
through the
DATA

6 OLAP browsers, pivot tables, vectors, measures—it's a whole different language. Jim explains it using a tool you probably have on your desktop. Excel.

BY JIM PAUL

DEPARTMENTS

- 3** From the Inside
- 20** Feedback
- 26** Newsmaker
- 28** New Products
- 47** Clif Notes

International Spectrum and MultiValue is a registered trademark of International Spectrum, Inc. All other registered trademarks and trademarks are the property of the respective trademark holders.

10 SOA WHAT? Still wondering what this SOA thing is? Mel explains the concepts of Service-Oriented Architecture in plain language. **BY MELVIN SORIANO**

16 PRACTICAL ESB Implementing Service-Oriented Architecture can be a daunting experience. Robert explains how Enterprise Services Bus assists the mid-market enterprises accomplish the goal. **BY ROBERT HOUBEN**

22 INVENTORY FOR PROGRAMMERS, PART 2 Charles continues his explanation of how to count widgets and why it is not always as simple as it seems. **BY CHARLES BAROUCH**

32 A BACKUP STRATEGY IN THE UNIDATA ENVIRONMENT Many shops have to take backups on the fly while users are still updating the database and just hope for the best if they have to recover from it. Jon shows us his solution to the “moving target” problem. **BY JON A. KRISTOFFERSON**

35 OVERVIEW OF MV.NET BlueFinity give us a tour of mv.NET and explains what is new with version 3.

36 AJAX AND OPENINSIGHT David takes us step-by-step in setting up a web page using AJAX with data served by OpenInsight. **BY DAVID GODDARD**

International Spectrum Regional Conferences

are coming up fast. I'm really excited about these upcoming events. We are sponsoring three regional conferences in the US, plus one international conference in Sydney, Australia:

Cincinnati, OH | **September 24 - 25, 2006**
Kings Island Resort

Seattle, WA | **October 11 - 12, 2006**
Holiday Inn Seattle SeaTac Airport

Sydney, AUS | **November 1 - 2, 2006**
Star City Hotel & Casino

New Jersey | **November 14 - 15, 2006**
Sheraton at Woodbridge Place

The regional conferences will be one of the first steps in implementing some of my vision for Spectrum's future. If you saw the article "New Leadership for International Spectrum" in the March/April 2006 issue, then you know that one of the many things I want to bring to the regional conferences is more training. With that in mind, I have worked out the following schedule:

1ST DAY

- 9 am - 12 pm • **Indepth Training**
- 12 pm - 1 pm • **Lunch**
- 1 pm - 5 pm • **50 Mins Vendor/Product Classes**
- 5 pm - 8 pm • **Table tops**

2ND DAY

- 9 am - 12 pm • **45 Min Vendor/Product Classes**
- 12 pm - 1 pm • **Lunch**
- 1 pm - 4:30 pm • **Indepth Training**

In addition to the in-depth training, I have asked the local MultiValue Users Groups to host their user group meetings/dinners on the 1st night between 6pm - 8pm. I believe the local MultiValue User Groups are very beneficial to the industry by providing a place for people to discuss their problems and find solutions that others may have already created.

As of this writing, I have the Seattle Pick User Group (SAPUG — www.sapug.org) signed up to host their meeting during the first evening of the Seattle Regional.

International Spectrum is also working with Open Systems Database Association (OSDA — www.osda.org) to help support their annual Exhibition and Conference in Florida this year.

If you are part of a user group in Cincinnati, OH or in New Jersey and would like to host a user group meeting during those shows, please feel free to contact me at nathan@intl-spectrum.com.

— NATHAN RECTOR, PRESIDENT, INTERNATIONAL SPECTRUM —
Nathan@intl-spectrum.com

NATHAN RECTOR
 President

CLIFTON OLIVER
 Managing Editor

TECHNICAL ADVISORY BOARD
 Henry Eggers, Independent Consultant
 Tim Holland, Holland Consulting
 Bruce Johnson, Tag Consulting
 Ronald H. Jones, Ron Jones Consulting
 Nathan Rector, Natec Systems
 Bryan Shumsky, Revelation Software
 Melvin M. Soriano, Eagle Rock Information Systems

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.



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 International Spectrum, Inc., 8956 Fox Drive, Suite 102; Thornton, CO 80260; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com. Copyright 2006 International Database Management Association, Inc. All rights reserved. Reproduction in whole or in part, without written permission, is prohibited.

PRINTED IN USA • ART AND DESIGN: CP Design; 619-246-6878; 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.

International Spectrum is a registered trademark and MultiValue is a trademark of International Spectrum, Inc. All other registered trademarks and trademarks are the property of the respective trademark holders.



GUI & Green Screen: Getting the Best of Both Worlds is Easy

Ever catch yourself thinking fondly of the past, when the MultiValue world was all green screen and users had achievable expectations of how their systems should look and feel? Little overwhelmed or confused about how you'll move those applications in a reasonable timeframe at an acceptable cost? Character To OpenInsight (CTO for short) from Revelation Software is the answer to your dreams.

CTO is a green screen application that runs on a GUI engine, allowing you to run your existing MultiValue applications without major changes to your code or data, while letting you gradually create the GUI components

against the same data at the same time. It allows you to continue to develop in a green screen app, or to make the GUI front end, or both. Heck, you can even call green screen windows from your GUI menus.

And cost? Revelation Software has the lowest prices of any of the MultiValue players.

For information about CTO, OpenInsight, or our Solution Partner Program, go to www.revelation.com/intl-spectrum/ishome or give us a call at 800-262-4747.

Make your dreams come true! Try the CTO from Revelation Software today.

CTO

CHARACTER TO OPENINSIGHT™

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.
45 St. Mary's Road
3rd Floor
Ealing
London, W5 5RG, UK
Phone: +44(0)208 912 1000
Fax: +44(0)208 912 1001
Email: info@revsoft.co.uk
Web: www.revsoft.co.uk

Revelation Software Australia Pty Ltd.
Suite 105
20 Dale Street
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

S O F T W A R E

Revelation Software is a division of Revelation Technologies, Inc. OpenInsight is a trademark of Revelation Technologies, Inc. All other brand and product names are trademarks of their respective holders.



BY JIM PAUL

Digging *through the* DATA

Understanding OLAP and Pivot Tables

Online Analytical Processing (OLAP) is a reporting and analysis technique that works with virtually any data. Its power lies in a “browser” (not to be confused with a web browser) that allows users to interactively create and change their own views of the data. A view is a tabular report and a corresponding graphical view that stay in synch.

Pivot Tables are an analysis tool built into Excel. A pivot table provides the same functionality as an OLAP browser. Microsoft has done an excellent job. Pivot tables are well thought out, easy to use, and arguably the best “OLAP” browser. An Excel pivot table can get its data from an OLAP server, such as OLAP Services that comes bundled with Sequel Server, or from a spreadsheet. This article is concerned only

with pivot tables created from a spreadsheet (but originating from data in a MultiValue database).

Pivot tables and OLAP databases are constructed of vectors and measures. Measures are things that can be counted and added such as charges, order dollars, number of appointments, or accounts receivable balance. Vectors (sometimes the word “dimension” is used) are everything else in the database — Vendors, Customers, Employees, Products, ZipCodes. Some vectors are derived from more granular measures, such as the way Product Line can be derived from an individual Product Code. All data has a special vector, time, which could be measured down to the minute for, say, appointment data, or as broadly as year, for some accounting data. *Continues on page 8*

INTRODUCING THE ALL NEW...

mitsreport

Advanced Report Generator



“Reports from over 1,000,000 records in seconds!”

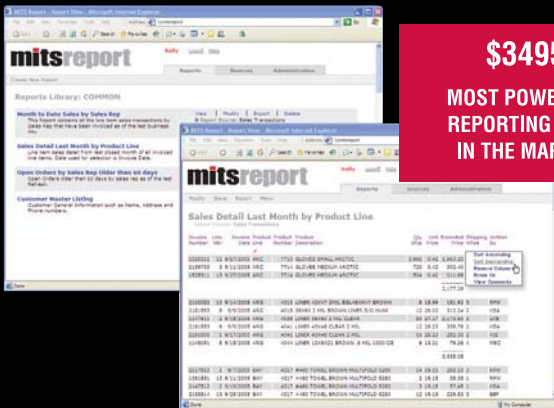
MITS Report has the power you have been waiting for in a report generator tool. Introducing an affordable next generation interactive tool from the industry leader in MultiValue database reporting.

FAST. Access, create and even manipulate reports from over a million records in seconds, not in minutes or hours.

SECURE. Protect any sensitive data with proven MITS based user or role access control technology.

EFFICIENT. Engineered for low burden on operational system even with multiple simultaneous users.

INTUITIVE. Easy to use for non-technical staff with no knowledge of the underlying file structure required. Just point and click for intuitive browser-based self-serve reporting.



\$3495
MOST POWERFUL REPORTING TOOL IN THE MARKET

Now you can have it all. While the Business Intelligence tool, MITS Discover, provides deep and dynamic data analysis to meet analytical reporting needs - MITS Report provides a simple, yet powerful tool for operational reporting. With MITS Report, users at all organizational levels can produce and restructure reports in a number of ways,

including modifying, adding, removing, or re-ordering columns without the repeated resource burden on your operational system—all with a query response time within seconds. Plus, with just the click of a mouse, data can be exported to Microsoft® Excel®, Adobe® PDF, and other popular file formats.

www.mits.com



Visit MITS today at www.mits.com or contact us at 888 700 6487 or info@mits.com

© 2006 Management Information Tools, Inc. All rights reserved. MITS, MITS Discover and MITS Report are trademarks or registered trademarks of Management Information Tools, Inc in the United States and/or other countries. All other company and/or product names may be trademarks or registered trademarks of their respective owners.

Understanding OLAP and Pivot Tables

Continued from page 9

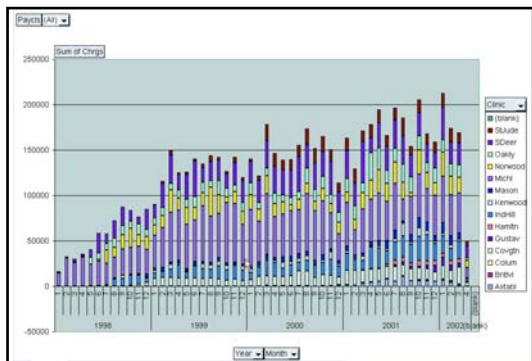


FIGURE 1

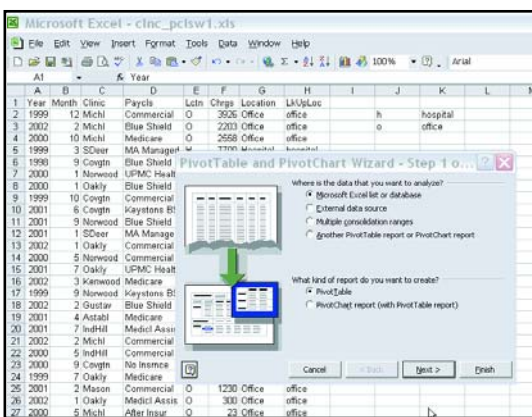


FIGURE 2

Figure 1 is a screen shot of a typical pivot table. The measure being charted is Charges. The values are plotted on the Y axis. The vectors Year and Month have been stacked and plotted on the X axis. The vector Clinic has been dropped to the right of the graph and results in the clinic codes being represented in different colors. The remaining vector, Location has been dragged above the chart. Vectors at the top of the chart are called “paging fields.”

The view in Figure 1 was constructed by a user starting with the spreadsheet in Figure 2. Each column heading becomes a “field” (Microsoft’s term for both vectors and measures) that can be used to construct a pivot table view. OLAP products from companies such as

Hyperion and Business Objects have browsers that look like Excel pivot tables but the source data is not a spreadsheet. It is a highly indexed, intermediate database called a “data cube.” Data cubes store summary data and may contain data from several sources, such as an order entry system and a job costing system.

OLAP is just one of many terms that refer to the same technology trend — Data Warehousing, Datamart, Business Intelligence, OLAP, Executive Decision Support Systems, and most recently, the word “enterprise” in some combination with the word “decision,” “analysis,” or “data.” The large number of names is due to new features being added as the technology matures, differences in emphasis, and marketers wanting to distinguish their offerings from their competitors.

Importance of OLAP

When managers are shown data represented as pivot tables or in an OLAP

FIGURE 3



cube browser, they are usually impressed. They see that the tool frees them from dependency on the IT department to extract data and allows them to interactively create tabular reports and graphs with little effort. They can quickly discern trends, weigh relevance, and gain insight into their operations. The tool is so powerful that access to the tool may even become jealously guarded. Garter Group recognizes the use of OLAP as one of four major trends. In a November 2005 report, they said,

“By 2010 IT departments in midsize and large companies will be 30 percent smaller than they are in 2005. IT jobs will be influenced by four major trends ... Business intelligence, online consumer services and collaboration will grow.”

Figure 3 shows the rapid growth of this market in dollar terms, even as the dollar cost of the products drops. A 2003 report from the same source, states:

“New analysis of data from The OLAP Survey 2 confirms that users of OLAP spreadsheet (Excel pivot tables) add-ins achieve more business success than any others.”

Excel Spreadsheets and pivot tables could very well replace all reporting. There is probably no report being done by your organization that would not be more useful if done as a spreadsheet with pivot tables. Excel is a product with which most business people are already familiar. Even if users are not familiar with pivot tables, they show far less reluctance to learn them than they do to special purpose OLAP products. Because companies already own Excel, the cost and risk of reporting projects based on it are far less than those of OLAP projects relying on specialized tools.

Creating a OLAP solution from your MultiValue Database

Microsoft describes pivot tables as using data from one of two sources—a spreadsheet or an OLAP server. There is a third way. You can use a MultiValue database as a repository of the same type of summary data stored in OLAP cubes. You can then load this data into Excel as a spreadsheet and conduct Pivot table analysis on it.

Programmers have long used intermediate files, often calling them summary files, to help with reporting. A carefully designed summary file can serve the same purpose as an OLAP data cube. A suitable summary file is created like this: Transaction files are processed and records in the summary file are updated. Vectors in the transaction file are used to form a compound item ID, while measures are accumulated in the attributes. For example, we could design a reporting file for medical billing transactions where the item ID is Doctor*Location*YearMonth and attribute 1 contains the dollar amount of charges. This file would be populated by a single pass through the transaction file. A number of reports could be produced from queries against this summary file.

A reporting file of this structure—a compound key of vectors with measures stored in the attributes—is ideal for pivot table analysis. What we need is a program to turn the records in the reporting file into rows in a spreadsheet. The program is simple—change the asterisks to tabs (Char(9)) in the ID and concatenate the attributes to it, again with tabs. Thus the first data row of the spreadsheet shown in Figure 2, was created from this multivalued record:

ID: 1999*12*Michl*Commercial*O

1. 3926 (charges in whole dollars)

All MultiValue systems have ways of transferring data to a spreadsheet. A

pivot table is created by highlighting all the columns in the spreadsheet, and running the pivot table wizard. Select “Data” and “Pivot Table & PivotTable” from the menu to start the wizard. Choose “Pivot Table and Pivot Chart,” and click “finish,” accepting all the defaults. You will be presented with a blank graph. The column headings become fields that you drag and drop to areas adjacent to the graph to construct whatever view you wish.

How to Boost your Career

Whether you are a consultant or an IT manager, enhancing your earnings and career depends on your ability to creatively deliver value to your customers and the visibility of your contributions.

Using pivot tables on a company's MultiValue database solves important business problems with clear thinking and simple implementation. If a company is already considering an OLAP solution to reporting, you will save them a heap of money by using tools they already own and are familiar with—their MultiValue database and Excel. If the OLAP concept is new to your customers, they will likely be impressed when you demonstrate pivot tables. After your first successful project, you may find managers from all areas of the company clamoring for projects of their own.

With thoughtful implementation, it is possible to address another perennial reporting problem. The problem of confidence in the data. OLAP projects address this problem in two ways: drill down and footing.

Footing means that the bottom lines of various reports from the same data source agree with each other. Thus, the

They see that the tool frees them from dependency on the IT department to extract data and allows them to interactively create tabular reports and graphs with little effort.

sum of charges by doctor for the month of June should agree with the sum of charges by location for the month of June. This is always true of pivot tables coming off the same spreadsheet and summary file. You will have even more confidence if the sum of the data in a spreadsheet can be reconciled to another source, such as the accounting system. Users may still question specific numbers on a report. Although it is possible that transactions were incorrectly classified, such as

crediting a procedure to the wrong doctor, transactions are definitely not lost.

Another common feature of OLAP products is the ability to drill down from high levels of summary to finer breakdowns. For example, the user may be suspicious of the total charges recorded in May for Doctor Smith. With a drill down feature, the user can see the individual transaction that made up Doctor Smith's May charges. Excel provides a drill down from the pivot table to rows in the source spreadsheet. But because the spreadsheet rows are themselves summary records, this drill down is not complete. It is up to the programmer to provide drill down from spreadsheet rows to actual transactions in the MultiValue database. There are simple and complex ways of providing this feature. The details won't be discussed here, but they depend on an index being maintained for entries in the summary file.

How to Begin at Your Company

Opportunities to introduce pivot table analysis abound. Look for suites of existing, related reports in your end-of-month process. Possibly a single summary spreadsheet can be substituted for all of them. *Continues on page 20*

SOA What?

Service-Oriented Architecture Moves into the Spotlight

BY MELVIN SORIANO

It's funny, really, how acronyms seem to be a given in technology today. We've become accustomed to taking these initials and making words and whole industries out of them. SOA (pronounced by many as "so ah") showed up a few years ago and now seems to be a popular one to discuss. At the recent Spectrum International Conference in Long Beach, I saw more than a few of us talking about SOA and its place in our market.

SOA stands for Service-Oriented Architecture. It's an array of many different services that communicate with each other. These individual services are independent and self-contained; they do not depend on the context or state of the other service. A key concept is "Information on Demand."

Services are similar to what we think of as subroutines. A lot of subroutines. A big difference is that SOA services work within a distributed systems architecture.

So, before I delve further, there is one thing I want to make clear. SOA is not SOAP. SOAP is an acronym for Simple Object Access Protocol. It's a standard for exchanging messages over a computer network, normally using HTTP and XML. It's really a messaging framework. It's just a minor portion of what SOA is all about.

Moreover, SOA is not a product, although many vendors now offer

products which can form the basis of a SOA.

No, Service-Oriented Architecture is more an IT strategy. It organizes and breaks apart the individual functions found in enterprise applications and then exposes them as interoperable, standards-based services. These services can then be combined and reused quickly to address ever-changing business needs.

Perhaps to better understand SOA, let's look at the problems this approach is meant to address.

Early neolithic programmers began to realize that writing business software was becoming increasingly complicated and nuanced. They needed a better, smarter way to reuse some of the code that they were rewriting.

We then began to program with modular design concepts, or subroutines. Programmers could write subroutines and functions and reuse their same code in different parts of the application. Many of us turned to 4-GLs to quickly handle the mundane processing for us. This was great for a while. It still is pretty good for many of us.

But a maintenance problem was developing. When a bug was discovered in a function somewhere, or if new functionality and interfaces had to be introduced, they had to track down all of the applications that used the function and

modify the code to reflect the changes. This first appeared to us in MultiValue with the introduction of GUI.

The industry moved to object-oriented programming and classes to solve this. But reusing code is not the same as reusing functionality, and complex extensive systems need to reuse that functionality. So IT moved to component-based software to handle its problems. Component-based software seems a good solution for reuse, but it fails to address the problems we have today. Technologically, we stare at complex issues like distributed software, disparate application integration, heterogeneous platforms, multitudes of protocols, ever changing devices, and of course the Internet.

SOA and web services provide a framework to address these issues. It may or may not stick around, just as the other approaches have come and gone. But by adopting a SOA, you can eliminate the headaches of protocol and platforms, and your applications may integrate more seamlessly. SOA builds upon the reusability of functions, so it may eliminate redundant data, rekeying of data, and that ever present human error problem. System design can focus more on strategic issues when these technical details are already handled.

Besides, today the main change that IT is currently facing is the shift to serv-

Informer Web Reporting for U2

<http://www.entrinsik.com/informer>



MultiValue Friendly, Intuitive Web Reporting. People Are Talking . . .

"... intuitive enough for technically inclined users to be able to set up their own reports without my IT staff."

- Lori Murray-Hawkins, University of New Brunswick

"Informer was a really quick way to get at data nobody could ever reach before. We found no other tool provided the graphical view into our multi-value database."

- Jordan McCall, Bratrud Middleton Insurance

"People in departments that we haven't introduced it to yet have seen other people using it and have approached us asking to get access to it."

- Joan Anderson, Gustavus Adolphus College

Important Features Include

Dynamic Table Joins
Direct Access via UniObjects; no need for ODBC
Drop reports to Excel, PDF, Email, or Saved-Lists

Role based security by function, data element
Scheduled report delivery to email or network
Friendly access to cryptic redundant U2 dictionaries



Entrinsik, Inc
919.848.4828
sales@entrinsik.com
<http://www.entrinsik.com>



SOA What?

Continued from page 10

ice orientation in general. Most of these moves to a service orientation are based on open standards-based computing. Exposing these services on the network and a service orientation hides the complexity of today's heterogeneous IT environments from business users.

So what is a service to us? A service is basically a subroutine that can be accessed via a standards-based interface. Each service represents a piece of coded functionality that maps directly to an explicit step in a business process. Services can be written from scratch or be reused as exposed modules of existing system applications. You should be able to see these services on the network.

In time, then, a catalog can be built up of the multitude of services. A robust catalog will allow business functionality to be accessed and reused across many different systems, in what one hopes to be a fluid and seamless manner. But, although SOA can support integration and consolidation activities within complex enterprise systems, it does not have a specific methodology for documenting capabilities or services. That's a weakness, but then, when isn't documentation a weakness?

In an SOA environment, the network makes resources available to other participants in that network as independent services. These services are accessed in standardized ways. Most implementations of SOA make use of web services, such as SOAP, in their implementation.

One can implement, however, SOA using almost any service-based technology.

Now, you're probably saying to yourself, haven't we all seen this before. True. Service-oriented architectures are not new. We've seen the Common Object Request Broker Architecture (CORBA) and the Distributed Component Object Model (DCOM), and they long have provided similar functionality. But these approaches were dependent on tightly coupled technology, not an open-systems approach.

So, unlike these traditional point-to-point architectures, SOAs are made of loosely coupled, highly interoperable services. These various services interact formally through a definition, or specifically a contract (a rather odd term admittedly). This definition is completely independent from the underlying platform and programming language of the services in question.

The interface definition therefore encapsulates (or rather, hides) the vendor and language-specific implementation.

MultiValue systems can be encapsulated, for example, and nobody would need to know what is happening on the back end.

Technologically, a SOA is completely independent of the development platform; write in Java or .NET or MultiValue and it shouldn't matter. The software components are highly reusable because the interface is based on standards and is independent from the underlying implementation of the service logic. A service written in MultiValue Basic could be used by a C# (pronounced C-Sharp) application. The reverse could also occur.

The SOA environment depends on four main areas: service, message, dynamic discovery, and web services.

Services

To recap, there are a few characteristics of services that SOA depends upon.

- The interface definition (contract) to the service is open and platform independent.
- Services can be located and invoked dynamically.
- Any individual service is self-contained, independent, and manages its own state.

Messages

Those that provide the services and those that use the services need to communicate, and they do so with messages. I know, rocket science. First, services expose their interface contract. This contract defines the what the service does, the messages the service can accept, and the messages the service can return. Since the contract is platform independent, the technology used to define messages must be open as well. Messages thus are typically written using XML documents that conform to an XML schema. XML describes all of the functionality, details, and scalability required by messages. These messages are crucial to the success of a SOA implementation; changing a message contract can be create havoc, so XML and its scalable messaging capabilities fits this need well.

Dynamic Discovery

At a high level, SOA is composed of three core pieces: service providers, service consumers, and a directory service. The directory is an intermediary between providers and consumers. Services register with the directory and consumers query the directory to find services. It's like a yellow pages to the SOA services.

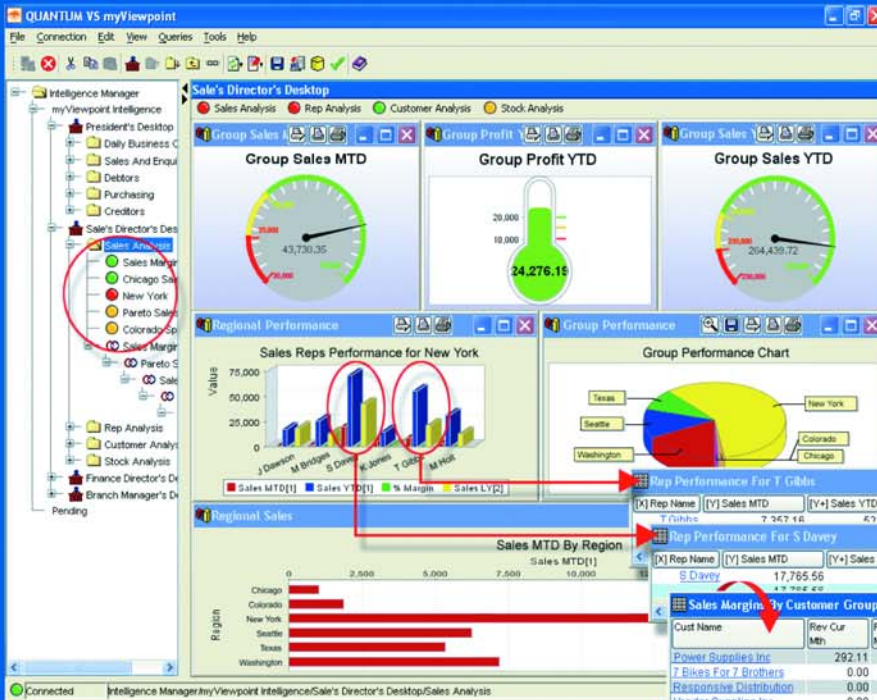
Web Services

Web services are a major component of SOA. They are built on top of well-known and platform independent protocols, which as one can guess, fits well

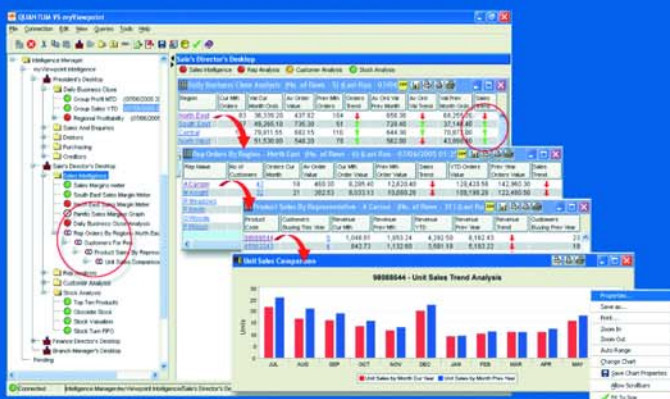
The main change that IT is currently facing is the shift to service orientation in general.

WARNING: Will Increase Profits

Transform raw corporate data into Business Intelligence
& maximize the value of your information assets.



- Personalized Business Intelligence compared against your Key Performance Metrics.
- Drill-down via Dashboard Graph Hotspots.
- Traffic Light Signaling Performance Alerts.
- Trend Analysis arrows.
- Hyperlink-anywhere & Drill-down through your Data Sets.



QUANTUM VS™ myViewpoint

The most powerful & cost-effective BI solution for the MultiValue DBMS User.

\$3,360.00 includes

- 1 x myViewpoint Professional &
- 1 x myViewpoint Standard.

Now available for UniVision, Universe, UniData and D3.

See us at the International Spectrum Conference, March 7-10, 2006. Long Beach, California. Via Systems Booth No.322 www.via.com

QM, A New Approach to MultiValue

Only \$119 per user
10 year upgrades just \$30
Proven for 4 years
Thousands in use
U2/Prime Compatibility
D3 extensions
Network Ready
Remote files
Client/server API
Useful bundled items:
AccuTerm + GUI
Coyote Web Server
Companion GUI/Web 4GLs
OpenTPH \$30/user
Neucleus \$72/user
Supported Platforms:
Windows 98, XP
Windows NT, 2000, 2003
BSD, Linux
OpenQM
GPL on Linux w/ source

EasyCo.com
800-470-2756

PICK PROFESSIONAL

Don't make the mistake of placing your career in the hands of just ANY search firm. We are professionals who understand the importance of CONFIDENTIALITY and RESPECT of a job search, and our database of clients is one of the largest in the country. Unlike the rest, we will work in YOUR best interests' to help you further your career. Because of our dedication and professionalism, we are recognized as the leaders in the PICK/UniVerse/Unidata placement industry in the Tri-State area and throughout the U.S. So if you are tired of putting yourself at the mercy of the rest.

CALL THE BEST! Contact...

Matt Hart

EXECU-SYS, LTD

1 Penn Plaza, Suite 3306
New York, NY 10119

(212) 967-0505

(800) 423-1964 x 302

Fax: (212)947-8593

Email: mh@eslny.com

Consultants Needed Nationwide

SOA What?

Continued from page 12

with the platform independent nature of the messaging and service delivery. These web service protocols include HTTP, XML, UDDI, WSDL, and SOAP. By mixing and matching these various protocols, you can pretty much handle all that SOA demands. You don't need these web services to do SOA, but they can be a huge assistance in meeting the open-system needs of SOA.

Will SOA be around forever? Client/server was once our main solution to everything, until the Internet made client/server solutions difficult to maintain and impossible to integrate. So, maybe not. There are problems with SOA that do need to be addressed.

One obvious challenge faced is managing the directory. SOA based environments can include many services which exchange messages to perform tasks. Depending on the design, a single application may generate millions of messages. Managing and providing information on how services interact is undoubtedly a fiendish, complicated task.

Security is another problem for SOA, as it is for most systems. Applications that use services, especially those beyond the firewall, are more visible and therefore more at risk to abuse.

Neither of these challenges seems too great to stop a good architectural framework, though. SOA is proving itself not just as an application architecture, but also as a business modeling philosophy. SOA is good at examining business services and operating models. SOA can provide a structure for IT to deliver the actual business requirements and adapt as businesses change. SOA as a business mapping tool can ensure that the new services represent the business view, not just what IT thinks the business services ought to be.

So, even though many of us have decades of experience in software development (am I that old?), we have yet to resolve the problems of software complexity. As complexity grows, as systems interact with other systems, as hardware and software interface options multiply, we have to find ever better answers. SOA and its use of web services is the latest answer.

The movement to SOA-based architectures will be evolutionary rather than revolutionary, always a good sign for long-term success. And although IT projects still are costly and take considerable effort, SOA approaches offer tremendous benefits. You get process orientation, adaptability, business application integration protection through abstraction, and cost-effectiveness. Now that's service with a smile.

Mel Soriano

(HTMMel aka XMMel) [IS](#)



MELVIN M. SORIANO

works at Eagle Rock Information Systems (ERIS), an

Internet Application Service Provider

and WebWizard/MultiValue Develop-

er. 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

www.eriscorp.com. You can always

call him directly at ERIS's Pasadena,

Calif., offices: (626) 535-9658.



FusionWare.net
Data Access Products Built on Standards

New! ADO.NET

New! Reporting Services

ODBC

JDBC

OLE DB

MV² SQL



**Gain Secure,
Easy Access to
Your MultiValue
Data from
Windows
or Java
Environments**



Future-proof Your MultiValue Systems

- Independent of database vendor or version
- Avoid forced version upgrades
- Easily integrate with non-MultiValue systems
- Extend MultiValue systems with non-proprietary reach

- Web Transactions
- SOA – Web Services
- Database Integration
- Reporting Solutions



Practical ESB for the Mid-Market

ESB technology makes it easier to build flexible services and to build security and policies into a SOA infrastructure.

“Enterprise Services Bus,” or ESB, is quickly becoming a mission critical, must-have for Enterprises, according to technology media and analyst communities. ESB is a technology that enables you to build a Service Oriented Architecture (SOA) interface within an enterprise; essentially, SOA can provide a means for users to access your existing proprietary applications via intranets and the Internet.

Media and analyst coverage of ESB technology tends to focus on applications of this technology within large enterprises. These firms are characterized as having numerous IT staff and large development budgets. However, there are hundreds of smaller sized, real-world enterprises who could benefit from ESB technology, but who must also contend with strained resources and lower risk on technology deployment. This is the class of organizations that truly comprise **Mid-Market Enterprises.**

Why Adopt ESB?

Mid-Market enterprises can realize many benefits by adopting ESB technology; here is a partial list:

It enables you to quickly build out services for your internal departments, customers, and partners without significantly increasing your ongoing costs.

It allows you to leverage existing staff and augment as you choose versus being compelled to hire an army of consultants.

Where outside consultants are needed to speed up deployment, you are able to hire developers who don't require highly specialized skills. (Moderate experience with

XML and Web Services is fairly common these days.) The results of what they build can be maintained by your current staff.

You can expect to see significant savings on all subsequent projects that you do, as you find that you can reuse DMZ services (middle-layer services running internal to your firewall).

You can take a build-as-you-need approach to deploying the services that are most critical knowing that these services are additive or can be easily be repurposed for other applications.

An ESB will allow you to create new functionality quickly in the immediate term and be able to take advantage of what you build for new applications as they become necessary. You should be able to do this with the skills that you presently have in-house with minimal training required.

An ESB should save you money and increase your agility immediately, over the option of coding the solution by hand, in the first project that you choose to work on. Over time the money that it saves you and agility that it gives you should increase as you build out a true SOA.

An ESB should also enable you to expose and automate systems designed to include numerous decision-making points by human operators. In short, your ESB should allow you to build links to your customers and partners in a cost effective fashion, without having to build out a whole new infrastructure or replace existing mission-critical systems.

What are the risks?

One of the biggest concerns that customers have about SOA is the security risk associated with anything that involves Internet technology. An ESB allows for a deployment to be fully secure, without requiring undue skills. Once having determined how to ensure a secure environment, ongoing maintenance of this security should be possible using your current staff.

The ESB that you choose should provide you with tools to enable you to mitigate common security risks.

Is Your Firm a Mid-Market ESB Candidate?

Mid-market enterprises have the following characteristics:

- Multiple and disparate enterprise-level systems at various divisions or locations, often acquired through mergers and acquisitions.
- Existing systems are mission critical, supported by knowledgeable but over-worked staff.
- Enterprises are usually not in "replace" mode with any of their systems. They are generally happy with what they have, but recognize the need and value of extending their systems to customers and partners in a way that requires them to expose these applications over the Web.
- Enterprises have close relationships with their solution providers and look to them to assist or even take the lead in recommending or providing solutions that enable access to new technology.

The IT departments within these mid-market organizations can be characterized as follows:

Continues on page 18

TIRED OF YOUR PRIMITIVE USER INTERFACE?

EH, PRETTY GOOEY, RIGHT?

NOT QUITE THE GUI I WAS TALKING ABOUT...

GO GUI WITH ACCUSOFT'S NEW GUI TOOLS! BRING YOUR PROGRAMS OUT OF THE STONE AGE TODAY!

TO LEARN MORE, VISIT WWW.ASENT.COM/IS TODAY!!




CREDIT CARD INTEGRATION

- Seamless Integrations into POS, Order Entry and e-Commerce
- Single or Multiple Merchants
- TCP/IP - Secure Internet
- Single or Multiple Modems
- Your Clearinghouse or Ours
- VAR Partnerships

Visa ■ MasterCard ■ AMEX
Discover ■ Diners ■ JCB
Debit Cards ■ Custom

TotaLink

1-866-796-7600
Sales@TotaLinkcc.com
www.TotaLinkcc.com



Get It Right From the Source.


Want control over the changes you make to your software?

PRC can help you get it right — and keep it secure.

For U2/MultiValue and SB+ environments

- ◆ Problem Reporting
- ◆ Change Control
- ◆ Project Management
- ◆ Time Tracking
- ◆ Test Management
- ◆ Deployment
- ◆ Rollback
- ◆ Audit

SJ+ Systems Associates
www.sjplus.com
954.796.9868



Practical ESB for the Mid-Market

Continued from page 17

- They generally consist of a combination of some internal people supplemented with consultants from their solution providers. This constitutes “their staff.”
- Their staff are business experts in the applications that they work with and are quite knowledgeable in the systems that these applications are built on.
- While aware of new technology, they typically do not have “expertise” in newer technology, such as Java, .NET, XML and Web Services.
- Their duties often span development, administration, and operations.
- They have little spare time to learn a new technology, let alone learn a whole array of related new technologies.
- They have limited budgets. While they can more easily buy software to solve problems, increasing permanent staff is typically MUCH harder and is generally not an option.

Setting ESB Architecture To Work

To put an ESB to work for your organization, you will need to draw upon SOA technology, because ESB and SOA work hand-in-hand to deliver applications to your users.

The concept of a SOA is actually not a new one, but with the advent of Web Services, and the many Web Services related standards currently evolving, the groundwork is being laid for users to rapidly and cost-effectively build connections between customers and partners. SOA also has the benefit of protecting partners and customers from changes in business procedures and underlying systems, providing both greater flexibility and greater stability.

An Enterprise Services Bus (ESB) is a technology that enables you to build a SOA interface in an Enterprise. In particular, ESB technology makes it easier to build flexible services and to build security and policies into a SOA infrastructure. There are many definitions of ESB, most of which are published by ESB vendors. The intention here is

to present a description of ESB that coincides with the requirements of most Mid-Market Enterprises.

Utilizing an ESB to build services in a layered manner brings immediate benefits to the development process. In particular, you can build a service layer that provides internal services, such as price-check, inventory update, posting to an accounting system, and others. Anytime that you have a sub process that you know will be useful to other people, it becomes a candidate for building out as a service.

If you think about a Purchase Order request service, there are a number of “services” that could be exposed by the underlying systems that would be useful to anyone within the Enterprise. These may not be something that you want to expose to the outside world without some security around them, but having them as services makes development of higher-level services much easier. It also makes your SOA less brittle, because these underlying services abstract out the various legacy systems whose functionality they expose.

Thinking about and building applications this way results in immediate savings as commonly needed components become immediately reusable.

What Does a “Service” Mean to a Mid-Market Enterprise?

A service is simply an abstraction of a high-level process.

What changes when a business process is performed within a SOA implementation is that all of the human components and decisions are replaced with automated components and business rules. The points in the process where a human is accessed are replaced with an integrated access to the related application systems. You may still have places where you want to involve a human agent, but ultimately the more you automate, the greater the value that you get from the SOA.

Some of the things that a human operator used to do that are now automated include:

- Identification: Something came in on the fax. Is it a PO?
- Transformation: On this customer’s PO, where do they put our account number for them?
- Integration: I need to access multiple systems.
- Workflow/Routing: As I access these systems, I need to make some decisions. I need some

“service-level” logic. I also need to follow predefined procedures.

- Communication: I need to be able to take inputs from numerous sources, and send responses to numerous targets.
- Security: I need to be able to ensure that customer XYZ can't access customer ABC's pricing and discount information, or check each other's order status.

FusionWare Integration Server— ESB for the Mid-Market

FusionWare Integration Server has been built from its inception to meet the unique needs of the Mid-Market Enterprise.

Benefits for the Mid-Market Enterprise

FusionWare Integration Server allows you to build out solutions in a layered manner. It allows you to integrate multiple mission critical or legacy systems, automate decisions around the data accessed from these systems, transform and route documents, and add workflow. All of this can be done securely, within your existing IT infrastruc-

ture, leveraging your current staff or your solution provider.

Benefits for the Solution Provider

As a solution provider, the relationship you have with your clients allows you to quickly recognize where a low cost, low risk integration solution would benefit them. In fact, they have probably approached you asking you if you know of an affordable way to extend their systems to their partners and customers.

In some cases you may have multiple customers with similar systems. With FusionWare, a solution that you build for one client can be easily resold to others. FusionWare delivers great value for your clients and repeatable, profitable offerings for your firm.

FusionWare is designed to showcase your core competencies and strengths without requiring a steep or lengthy ramp-up process to learn the technology. Your staff


can get up to speed within hours, and be ready to create and deploy tested services for your clients.

The FusionWare Integration Server reduces the financial and overall business risk associated with “not to exceed” or “fixed bid” constraints in today's competitive market. FusionWare allows you to profitably deliver integration projects on time and on budget.

While providing ESB and SOA capability to many databases and disparate systems, including SQL Server, DB2 and Oracle, we are very familiar with MultiValue databases, which gives us a clear understanding of the issues facing you. Our ultimate goal is to provide you with the best options to address your changing business needs, while retaining the MultiValue applications and business logic that have been running your business for years. With our core MultiValue data access products or the FusionWare Integration Server we enable integration in new and powerful ways. **IS**

CONTACT info@fusionware.net to learn how FusionWare can help you reach your business performance goals faster.

Looking for someone who really knows how to handle RedBack? and now .NET!!!



Come to AdvancedWare The .NET & RedBack Experts

Our Real-Time Web and Wireless PDA solutions work with any UniVerse or UniData application.

Internet/Intranet Solutions
Wireless PDA Applications
RedBack Training & Development

As the world leader, AdvancedWare has more successful Web and PDA implementations than anyone else. Give us a call to unlock the power of your UniVerse or UniData system.

949-609-1240
www.AdvancedWare.com



Understanding OLAP and Pivot Tables

Continued from page 9

Listen for complaints about the accuracy of existing reports. Any request for a new report is an opportunity to use pivot tables. Any quality oriented initiative at the company will rely on business metrics. Total Quality Management, Continuous Quality Improvement, Six Sigma, Management by Objectives, ISO 9000 projects depend upon measurement well served by pivot tables.

Once a user is introduced to pivot tables as an alternative to paper reports, they will likely request that all future reports be done as spreadsheets and pivot tables. As programmers become familiar with the technique, they will find it faster to produce a spreadsheet than fool around with the query language, correlatives, and the limitations of paper.

As you move in this direction, you will soon encounter complex reporting requests involving data from multiple

sources. You will find it necessary to come up with enterprise wide definitions of vectors. For example, there may be a need to combine data from appointment scheduling and billing. Appointments are made with doctors on particular days, and the billing system contains the actual charges by a doctor on an individual day. But the appointment system may not use the Doctor file used by the billing system. If data is to be combined from these two sources, doctors in the appointment system will have to be mapped to doctors in the billing system. The OLAP term for the importing and manipulation of data prior to analysis is called Extract, Transform, and Load (ETL).

OLAP products that do not use pivot tables and have their own OLAP database do have a place. Because of the size limitations on an Excel spreadsheet, you will reach limits in the number of vectors you can use in a single

pivot table more quickly than you will with a third party product. However in the areas of cost, speed, and user acceptance, you will find Excel the best choice. If you are a software vendor, you will find yourself in the enviable position of having created a desired enhancement to your core product, one that is perfectly integrated with your software, and which costs you nothing to re-sell. is



Principle in Northtec Consulting Group, which he formed in 1985, Jim Paul speaks at dozens of conferences around the country

and Europe and has written articles for several different publications, including Computer World. He has had clients and distributors in Europe, Australia, Russia, South America, and the Caribbean and has written complete software applications for finance companies, manufacturers, HMOs, and medical providers. A Sun certified Java programmer with a BA from Indiana University of Pennsylvania, he has taught college level courses in math, accounting, history, and economics.

He can be contacted at jimpaul@fuse.net.

Feedback

What came first, the letters or the letters-to-the-editor department?

We are making another change to International Spectrum Magazine—a Feedback Department, sometimes known as Letters to the Editor.

We want to hear your comments, your reactions, your agreement or disagreement with what you see. Also, do not hesitate to let us know about things happening in the MultiValue Community we may not have heard about yet.

Please send your comments by e-mail to: editor@intl-spectrum.com

REGISTER TODAY!

Bringing the International MultiValue Community Together



New this Year!

- ◆ **The New Generation** of Web-Based Application Generators for MultiValue
- ◆ **Corporate Updates** From the Core MultiValue Database Providers
- ◆ **The Latest** on MultiValue and Microsoft's .NET
- ◆ **The Newest** in MultiValue Business Intelligence Tools
- ◆ **Using Your** MultiValue Data with Microsoft's Word and Excel
- ◆ **15 Of the Newest** Software Tools Including OpenQM for MultiValue



Cincinnati, OH

September 24 - 25, 2006

*Kings Island
Resort*

Seattle, WA

October 11 - 12, 2006

*Holiday Inn Seattle
SeaTac Airport*

Sydney, AUS

November 1 - 2, 2006

*Star City Hotel
& Casino*

New Jersey

November 14 - 15, 2006

*Sheraton at
Woodbridge Place*

REGISTER TODAY!

www.intl-spectrum.com

Inventory *for* **(Part 2)** Programmers

Instinctively, we want to see things simply. As a result, many systems try to treat inventory as a single concept: If I have it, it's inventory. Unfortunately, the world doesn't comply with the simple view. Raw Materials, Finished Goods, and Imports all have to be dealt with differently, even though they are all kept in inventory.

Why are Imports Different

From a manufacturing standpoint, it takes as many yards of Italian silk to make a scarf as it would take of Chinese silk, or American silk. We still can't think of imported Raw Materials as Raw Materials, because on the one hand our respective governments don't see them as the same, and on the other hand, our customers don't see them as the same. If we make a scarf that bears the label "Pure Japanese Silk", it had better be true. Just like we don't lump all the colors of fabric together, we can't lump all the countries of origin together. From a governmental standpoint, the receiving, use, export, and sale of foreign materials

may have tariffs and quotas associated. This is a case where not tracking country of origin in your software might be a federal offense.

Additionally, some categories of Imports require documentation to travel with them. Moving to the Finished Goods side of Import, we have to consider that cars require Certificate Of Origin paperwork, which must be associated with the car from your receipt of it to the moment you sell it. Clearly, Imports have to be treated separately than domestic goods, regardless of which country you call home. This goes doubly if you are going to Export the goods at sale time, or as a transfer between your own offices and facilities.

Why Raw Materials Aren't Finished Goods, Except When They Are

Accepting that Imports require special handling, we still have to ask why inventory systems need to segregate Finished Goods and Raw Materials. While it is obvious that manufacturing needs to see them as different, the inventory question is hazier. Let's start by defining the terms.

Raw Materials are all the non-labor components that I can add to a Bill of Materials in order to create Finished Goods.

Finished Goods are anything you can sell as-is.

Obviously, there's an overlap. For example: If you sell coffee, then bags of

coffee are a Finished Good, because you can sell them. However, if you sell a gift basket which contains three bags of coffee, then coffee bags are a Raw Material (part of the Bill of Materials) needed to make a gift basket.

So, if they overlap, why bother thinking of them separately?

Location, Location, Location

A good system always minimizes the chances that a user will be able to easily do the wrong thing. Storing raw materials near shipping and far from manufacturing is a problem. Your software can't watch for this problem if you don't know which inventory items are in each (or both) of the categories. Likewise, when your software allows the construction of a Bill of Materials, we shouldn't allow every Finished Good to be used as a Raw Material, just because some of them can be. In our example, it is unlikely that the resulting gift basket is a Raw Material for some other Finished Good. It is likely that a confused user might, when building a Bill of Materials for a gift basket, put the Finished Good ID in the wrong place. We can't catch them if we don't track the role that each item in inventory plays.

We also need to think about how we report inventory. While the plant manager might need a total inventory report for the physical location, most people who care about stock only care about one type. That requires us to know which is which.

We also need to think about how goods are taken out of inventory. Production reduces Raw Materials, Sale reduces Finished Goods. Both are reduced by shrinkage. When something atypical comes up, like Production seeming to consume Finished Goods, we need to know it's unusual so that we can test the assumption programmatically or report it for review. The more we know about an item in inventory, the better we can find fraud, errors, theft, and other profit killing behaviors.

It's Not Mine, I'm Holding It For a Friend


Obviously, these three forms of inventory do not constitute the entire list of possible items. There are other types of inventory, like third party inventory. One of my customers has a system for pass through inventory that allows his customers to sell goods directly out of his master inventory. The variations are endless.

Prejudice

One final note: when you change vertical industries, rethink everything you know about inventory. At a certain level, every business handles inventory the same way, but the small twists can loom large in deciding how your software should treat their system. I, for example, had a clothing manufacturing background as my first exposure to inventory. I started learning inventory for clothing when I was six and visited my father at work. I'd climb into the bins and toss the robes to the shipping clerks. It is easy for me to see every inventory system as if it were a garment inventory.

Doing that would stop me from understanding the fine points of other people's system. I've even run into garment businesses with very different inventory models than my dad's. Some examples:

1 I was used to storing in dozens. One company I worked with stores in units because they have different quantity arrangements (two packs, three packs, five packs) with each customer. They never ship a prepacked dozen, ever.



GREAT!

Celebrating 8 years as a leader in the Multivalue Industry

www.pickprogram.com
(614) 921-9840

Consulting Services

PPS delivers high quality solutions to the Multivalue Industry

Our mission is to resolve our clients' requests by creating effective applications in a cost efficient manner

Up Front Estimating (know your costs before starting the project)

Any size project (Hourly, or project based)

Local and Nationwide Support

Our consultants average 12 years MV experience

Do you need Multivalue questions answered? Contact us for a quick response. Initial contacts are always FREE

Technical Expertise

All Multivalue platforms (Universe, Unidata, D3...)

Web Services

Open Database Connectivity

.Net and Visual Basic to Multivalue programming

Microsoft Office integration

MITs development

Project outsourcing & management

System migration

Technical support and training

Barcode application programming, hardware, supplies, sales and support

Onsite and offsite services

Barcode Services

As a logical extension of our consulting experience in barcode applications, we now sell and service:

Barcode Design Software

Barcode Labels

Card stock

Custom Forms

Preprinted Labels





Printers and Scanners

Custom software programming for all barcode symbologies.

Best price for the value

Over 50 years in barcode expertise

Free quote - tell us your current specs, volume and pricing - receive a quote within one week

OUR PRICES ARE GREAT!... OUR SERVICE IS GREAT!... OUR SUPPORT IS GREAT!.

2 I was used to systems where each color was a separate order detail line. One company shipped preset mixed color packs only. You couldn't order six of the red ones. Storing in single color dozens would never be efficient. We stored in finished prepacks.

3 I was used to systems where backorders were preferable to delays. One company I worked with had ship complete agreements where they shipped all or canceled the whole order for certain customers. I had to rethink when I allocated and deallocated inventory.

Obviously, anyone who has been in this industry long enough could add their own twists to this list. Once we

think of inventory a step in the chain, we look at how we store things based on where they come from, how they are used, how they are made, and how they are shipped. It's far from intuitive, no matter how appealing the "I count six, I have six, I can sell six before I run out" might be.

Part III will discuss the importance of tracking the inventory you don't have. is

CHARLES BAROUCH is the CTO for Key Ally, Inc. and Mount Olympus Systems, Inc. He is also President of the U2UG and a regular Spectrum Magazine contributor.

Charles Barouch Results@KeyAlly.com - Consulting
Zeus@MtOlympus.us - ETL/Migration/Integration
(718) 762-3884x1



DesignBais

Rapid Application Development Environment

Our search has ended. In my 25 years in this business I have not seen a development tool that fits our industry as well as DesignBais. It simply does everything necessary to modernize our application without overloading us with options we do not need.

In two weeks we were able to train all of our programmers and get a significant portion of our application converted. I was amazed at how intuitive it was for our Pick programmers.

I am convinced this will become the de-facto standard for MultiValue applications moving to a browser-based interface. It is really a world-class product that is ready to place the thousands of MV applications well ahead of their competition.

For us at Quality Software Systems, Inc., it will not only rejuvenate and enhance our Laboratory Information System but allow us to enter new peripheral markets.

Sincerely,

Jeff Caspari
President - Quality Software Systems, Inc.

- High Performance
- Robust and Reliable - Release 3.3
- Extensive functionality
- Multi-byte capable
- Utilizes existing resources
- Affordable
- Multi-platform Support
- Complete application development environment

Download an evaluation copy today

Multi-value has never looked this good!



www.designbais.com

DesignBais International names Nebula R&D for Professional Services



DesignBais International has partnered with Nebula Research and Development to provide extended Support and Development services for a growing base of international users and developers. As part of this agreement, Nebula R&D will assist developers, process requests for support and enhancements, and otherwise serve as a backup resource to the existing international client services staff for the flagship product DesignBais.

David McLean, Director of DesignBais international believes that Nebula brings a wealth of experience and expertise to the DesignBais team. "Tony Gravagno, the President of Nebula R&D, has been a staunch supporter of the DesignBais product since the first release over two years ago. I am thrilled to have Nebula R&D as a partner as we continue to grow internationally."

DesignBais offers in-depth training and professional services, which provides a serious kick-start for new developers.

"We want to ensure that our training encompasses the developer's

own applications, not just another Widgets example", McLean said. "When they finish a training course, they leave with functional components of their own application, which has proven to be an extremely successful outcome."

This new relationship with Nebula allows DesignBais International to extend the level of professional services and training available, particularly in North America. Nebula R&D will also provide pre-sales assistance, and custom development services through DesignBais International and as an independent consulting agency.

This new relationship with Nebula allows DesignBais International to extend the level of professional services and training available, particularly in North America.

"I've had nothing but praise for the DesignBais product since I first saw it," says Tony Gravagno. "More importantly, VARs and end-users are going production with attractive and highly functional web interfaces, and their testimonials reflect their satisfaction with their vendor, the process, and the results. We're proud to have this new relationship to work directly with DesignBais International, the DesignBais product, and with their client base."

DesignBais is a functionally rich toolset that allows MultiValue developers to design and create enterprise wide web-based appli-

cations without the need for learning new technologies. The resulting applications are deployed at many sites throughout the world. DesignBais is compatible with the major MultiValue database implementations. DesignBais International has offices in North America and Australia. The product is sold and supported by a global network of value-added resellers and distributors. Visit them at www.DesignBais.com. ■

Nebula Research and Development is a recognized provider of development tools and services for a worldwide marketplace of MultiValue/PICK and Microsoft .NET users.
<http://Nebula-RnD.com>



Placentia-Based Computer Firm takes on Project for the New York Philharmonic

The New York Philharmonic has hired Placentia-based SRP Computer Solutions, Inc. to upgrade its Revelation based computer software system with utilities designed to bring the world-class orchestra into the 21st Century.

The programming will enable the orchestra to access payroll, scheduling and other applications from remote locations, and it integrates all components into one easy-to-navigate system.

"Our original DOS-based program was becoming unstable," said Elizabeth Lee, Information Systems

manager for the New York Philharmonic. "SRP's FrameWorks program allows us to modify and create functions under a Windows XP format."

SRP helped develop a system that accommodates all payroll needs, allows for regular maintenance of a library network dating back to 1842, and establishes a more flexible scheduling system. It can also be used remotely with multiple levels of security – something other program upgrades couldn't offer, she said.

Lee said the Philharmonic hopes to create a web-based version of the program to allow the public access to library data, schedules and more. Eventually, people may be able to log on to the New York Philharmonic web site and see what works were being played in concert more than 150 year ago.

"(SRP) was very flexible with us," Lee said. "They worked with us to give us what we need."

SRP Computer Solutions, Inc. is a leading developer of OpenInsight and Advanced Revelation utilities. They also provide consulting, programming, and training services for all Revelation products. The corporation has served a variety of clients throughout the U.S., from hospitals to manufacturers, horse breeders to shipping companies. ■

For more information contact:

Don Bakke, President, SRP Computer Solutions, Inc.
(714) 577-9317

OpenInsight Finalist At LinuxWorld's Product Excellence Award

Revelation Software is proud to announce that their flagship product, OpenInsight, was selected as a finalist in the Best Application Development Platform or Tool category at LinuxWorld Sydney's Product Excellence Awards.

"OpenInsight is the mature application development environment that Linux developers have been looking for," said David Goddard, Managing Director of Revelation Software, Australia. "This award confirms that OpenInsight is a market leader in application development tools for both Linux and Windows platforms."

In announcing the award, Revelation Software was complimented on their contribution to the Expo and the Open Source Industry. ■

Ladybridge Systems Ltd Announces Support for Object-Oriented Programming in OpenQM



The OpenQM MultiValue database environment now natively supports object-oriented programming techniques and structures. Ever since MultiValue Basic was first deployed on early MultiValue systems, programmers have struggled with managing larger and larger projects. Existing programming techniques like modular programming design and variable scoping, while useful for small applications, become less effective as applications get larger and more complex. By adding true object-oriented run-time support to QMBasic, OpenQM becomes the first production MultiValue environment to add native object-oriented functions in it's core application environment.

Object-oriented programming is all about code maintainability and reuse. The first technique most programmers employ is to use external subroutines. While

subroutines are effective, they still lack the full level of code and data abstraction and reuse provided by objects. With objects, MultiValue application developers can build reusable application logic with power and ease previously unavailable in any MultiValue database environment.

OpenQM objects are true to their MultiValue roots. All objects are typeless, fully scoped, nestable, and run-time resolved. By including support for run-time properties with parameters as both rvalues and lvalues, OpenQM objects create a natural, extensible development environment that optimizes programmer productivity. OpenQM objects are also a native part of the QMBasic language and runtime. This ensures seamless integration with existing applications and excellent run-time performance.

The addition of objects to QMBasic brings modern object-

oriented programming techniques to MultiValue applications. Object-oriented programming is widely taught and is supported in most modern application development languages such as C++, Java, and Perl. The addition of object-oriented support in OpenQM will give MultiValue developers the code reuse and modularity tools needed to compete with other development environments. ■

For more information about OpenQM contact Ladybridge Systems or visit www.openqm.com.

Ladybridge Systems
17b Coldstream Lane,
Hardingstone, Northampton,
NN4 6DB, England
www.ladybridge.com
www.openqm.com
sales@ladybridge.com

New Products

DesignBais adds BlueFinity's mv.NET to Product Offering



DesignBais International and BlueFinity International are pleased to announce a worldwide agreement under which DesignBais will market and sell BlueFinity's mv.NET as an option to the DesignBais development environment.

mv.NET is a suite of components designed to provide developers with a comprehensive solution to the challenge of creating Microsoft .NET based applications requiring access to Multi-Value databases. With support for all major MultiValue platforms including D3 and jBASE among others, mv.NET provides an alternative access mechanism for DesignBais applications.

"The timing for this relationship could not be better," comments David McLean, Director of DesignBais, "with our recent growth we are facing increasing demand for DesignBais on a wider variety of MultiValue platforms. mv.NET is an attractive option for DesignBais customers which eliminates extra software layers and streamlines set-up."

DesignBais is a functionally rich development environment that enables MultiValue developers to design and create browser based applications. It was specifically designed to be compatible with MultiValue and multi-dimensional databases. At DesignBais' most recent release of 3.3, mv.NET is supported on D3 and jBASE, with additional platforms to follow.

"mv.NET works with all major MultiValue platforms. There are no other .NET development products with this scope," states Pete Loveless, CEO of BlueFinity International. "This partnership will allow DesignBais to open up their product to the whole of the MultiValue community; we look forward to joining them in delivering the mv.NET advantage to their customers."

BlueFinity International supplies leading-edge .NET oriented software development tools and consultancy services to the MultiValue database and Microsoft developer communities.

About DesignBais International

DesignBais is a functionally rich toolset that allows MultiValue developers to design and create enterprise wide web-based applications without the need for learning new technologies. The resulting applications are deployed at many sites throughout the world. DesignBais is compatible with the major MultiValue database implementations. DesignBais International has offices in North America and Australia. The product is sold and supported by a global network of value-added resellers and distributors. Visit them at www.DesignBais.com.

About BlueFinity International

Founded in 2002, BlueFinity International offers its two flagship products – mv.NET and RSDC – to the global MultiValue community. mv.NET allows the full benefits of the .NET service oriented architecture technology to be realized by users of established MultiValue applications; RSDC enables MultiValue developers to utilize the very latest Microsoft Reporting Services technology.

Visit www.BlueFinity.com for more information.

Contact: Info@DesignBais.com; +1-508-686-2817



BlueFinity Unveils mv.NET Version 3

BlueFinity International, a member of the Mpower1 group of companies, is pleased to announce the general availability of mv.NET Version 3. mv.NET is a powerful and extensive toolset which enables MultiValue developers to harness the full power of both their MultiValue system and Microsoft's .NET application platform. Its unique set of features allows the developer to use the latest Microsoft technology while still leveraging all of the traditional benefits of MultiValue functionality.

The mv.NET architecture is designed with both performance and flexibility in mind. This, combined with an implementation that provides seamless integration with the Microsoft Visual Studio .NET environment, provides a potent Rapid Application Development environment fine-tuned for the MultiValue developer.

Version 3 is the start of a series of releases aimed at introducing a range of additional application development productivity tools built upon the solid, performant foundation that mv.NET already possesses.

"mv.NET Version 3 is the starting point of a major investment and enhancement program for mv.NET that will add further to the range of application development productivity tools already provided and will maintain mv.NET's position as the .NET development toolset of choice of MultiValue application developers around the world," states Ian Leeper, Director of Mpower1.

A major addition to the growing list of enhancements available in this latest version is Binding Objects support for Web applications (webBO). All of the advantages of mv.NET Binding Objects technology have been extended to the development of browser based applications with the introduction of full support for WebForms. This brings all the RAD benefits of faster time to market and easily supported code to web application developments as well as for Rich Client.

"Binding Objects provides the necessary functionality for a modern development organization to deploy applications that will compete head on with anything seen by the industry and to fully exploit the key advantages of the underlying MultiValue database," explains David Cooper, Lead Developer at BlueFinity International.

Included within the mv.NET suite of components is a sophisticated ADO.NET managed data provider, known as the Adapter Objects package of components. This provides a robust, high-performance data retrieval and update capability to developers who prefer to use ADO.NET as their data access methodology.

Adapter Objects provides a 100% ADO.NET compliant, pure managed code data provider. This industry standard technology is then augmented by a range of sophisticated extensions which allows the full power of MultiValue systems to be harnessed from within the ADO.NET framework.

Adapter Objects supports all of the links to other associated and commonly used .NET classes and components, such as typed datasets and databinding. It is also fully integrated with the Visual Studio IDE environment.

Cooper explains, "Uniquely with mv.NET, you get a broad range of components, each providing a slightly different slant on the challenge of providing .NET to MultiValue interoperability. If ADO.NET does not provide the power or flexibility demanded by your application, the other mv.NET technologies can be used to complement, where necessary, your use of ADO.NET."

BlueFinity also recently announced the availability of its new Reporting Services Data Connector product (RSDC) which, when used with mv.NET, allows SQL Server Reporting Services (SSRS) users to utilize MultiValue data sources in their SSRS authored reports. By utilizing BlueFinity's mv.NET as the database connectivity solution, the RSDC product is able to offer unparalleled SSRS integration across a wide range of MultiValue platforms.

"This is the start of a very exciting period for BlueFinity," says Leeper. "We have some great visions as to how we can help MultiValue developers perform more competitively and more productively in the arena of application software creation. It is our aim that during 2006, the MultiValue developer community will recognize mv.NET to be the leading product in its class and that it can provide them with significant competitive advantages in the creation of .NET applications."

mv.NET and RSDC are currently available on D3, jBASE, mvBASE, Reality, Power95, QM, UniData, UniVerse, UniVision and others by request.

For a full list of Version 3 enhancements or to find out how to request a free thirty day evaluation, visit www.bluefinity.com.



Sumo Computer Announces Low Cost MultiValue Platform

Sumo Computer was created on the premise that low cost Open Source application platforms will be the servers of the future. With the vast base of Open Source applications available, the problem is not the availability or cost of the Open Source systems but the time it takes to configure the Open Source applications to work with one another. Sumo Computer provides that configuration service on a low cost, high performance platform.

Sumo Computer uses the Buffalo Technology Kuro Box/HG as it's low cost platform. The Kuro Box/HG is a small headless server measuring 2.4" x 7.0" x 7.4". Sumo Computer adds a Maxtor 250 Gb disk drive and Gentoo Linux. The Kuro Box/HG sports a 266 Mhz 603e PPC processor with 128 Mb of memory. The Kuro Box/HG also has a Gigabit Ethernet interface and two USB 2.0 compliant interfaces.

The Sumo Computer configured Kuro Box/HG costs \$399.00 USD. This fits within the low cost parameters that were originally targeted by Sumo Computer. While the performance of a 266 Mhz 603e PPC processor is not as fast as newer chips, it is still fast enough to perform many tasks in a satisfactory manner. The processor and speed are roughly equivalent to an IBM RS/6000 model F50. The configuration of the Kuro Box/HG is

fixed and cannot be upgraded without user modification of the motherboard.

The different models that Sumo Computer configures are all based upon the same hardware platform while offering different software configurations based upon popular Open Source packages. Many software vendors are now using the Open Source paradigm as a marketing tool for their applications. They offer their commercial applications as Open Source applications. If you wish to acquire licenses and software support the software vendors will sell you their commercial application.

The current three models provided by Sumo Computer are the OpenQM Kuro Box, the LAMP Kuro Box/HG and the Configured Kuro Box I. These are the same hardware platform with Gentoo Linux and a 250 Gb disk drive, but with subtle variations in the Linux configuration.

The **OpenQM** model has the Open Source version of Ladybridge's QM Multivalued database system installed. The LAMP model has Linux, Apache2, MySQL and Perl/Python/PHP installed. And finally, the Configured Kuro Box I has Gentoo Linux installed with some general Linux utilities to make the administration of the server easier.

Gentoo Linux was chosen for it's vast store of PPC applications, ease of administration and even easier package installation features. To install one of the over 10,000 applications ready for Gentoo Linux you only need to use the Gentoo "emerge" command to download and install it. You also use the "emerge" command to keep your version of Gentoo Linux up to date.

Continues on page 30

New Products

Continued from page 29

Some example applications that the Sumo Computer configured Kuro Box/HG servers excel at are:

- ◆ OpenQM application server
- ◆ Rsync servers that can be used to backup daily snapshots
- ◆ DHCP and DNS servers
- ◆ Apache2 Web Server
- ◆ LAMP server
- ◆ Samba Server to tie Main-frame applications to Windows Desktops
- ◆ Multimedia server for MP3 or Video applications
- ◆ VoIP server

In all, a small, low cost, high performance server pre-configured with Open Source applications and a MultiValue environment for testing and development is a lofty goal. The Sumo Computer configured Kuro Box/HG system provides such a system.

For more information, visit our web site at www.sumocomputer.com.



mits

MITS Discover 6.0 Launches with 30+ New Features

New Release of the Leading OLAP Business Intelligence System Focuses on Usability Enhancements

MITS is announcing the release of a new version of MITS Discover, the OLAP Business Intelligence system for interactive analytics.

MITS Discover 6.0 is currently available for IBM UniVerse and UniData, with support for other MultiValue databases to follow. Packed with more than 30 new features, this new version is offered at no cost to users on the MITS annual support plan.

The move from MITS Discover 5.4 to MITS Discover 6.0 shows the company's deep commitment to continue improving productivity for the more than 3,000 end-user organizations, resellers, and systems integrators currently using the software.

Vast Usability Improvements

The release of MITS Discover 6.0 marks the most significant usability improvement in the history of the product. This substantial enhancement focuses on MITS Discover's popular Web-based user interface, through which users can explore MITS Discover applications in their preferred Web browsers. Current MITS Discover users will appreciate the next-generation user interface in version 6.0, and new users will find its intuitive user interface easy to navigate.

"Over the past six months, we have committed our resources to usability testing and meeting with our customers and partners to assure substantial usability improvements in version 6.0," Fred Owen, president of MITS, states. "I believe this release represents a strong step forward in the power of the product."

A highlight of this MITS Discover release is the new MitsWeb interface, which has an updated look and feel. New drop-down menus replace pop-up windows, giving users fast access to streamlined data summaries and drill down. In addition, now it's easier for users to add or delete

columns, print entire flash screens (by exporting to Adobe PDF), print currently displayed rows of flash screens, and export to Microsoft Excel.

MITS Discover 6.0 offers an expanded security model that includes a password expiration feature and the ability to restrict access to information on every level.

Better Data Visualization

In MITS Discover, dashboards give users a way to track their companies' key metrics at a glance, with information pulled right from the application. In response to direct requests from customers and partners, the company has improved dashboard functionality in version 6.0 of the product. Dashboard improvements include: an updated, Stoplight dashboard object that changes color based on the data; the ability to show, hide, and resize the dashboard; automatic grouping of dashboard areas based on headings; and one-click access to edit the dashboard. These enhancements represent the company's intention to provide better data visualization as MITS Discover continues to mature.

Enhanced Detail Display

The list of detailed transactions no longer appears as a static view that users must page through to get to the data they need. Detail display in MITS Discover 6.0 offers the following updated functionality:

- The ability to sort by columns and set the default sort order for detail display, with up to four different levels of sorting
- The use of hyperlinks in columns to allow users to inte-

grate MITS Discover with their outside applications

New Architecture for Browser Accessibility

The Java Gateway in MITS Discover 6.0 gives users an alternative to the existing Windows-based gateway. While the older gateway is still available for MITS Discover, the Java Gateway offers greater platform-independent support and improved charting. The release of the Java Gateway delivers a new Charting Engine, so users can choose from a variety of chart types, including two new dashboard objects (Speedometer and Thermometer) and eight of the most popular charting options, including Line Charts, Stacked Bar Charts, and Area Charts.

For more information on MITS Discover 6.0 or to see a live demo of the product, contact MITS at 888-700-6487 or info@mits.com.

About MITS

MITS is the leading developer of Advanced Reporting and Business Intelligence for the Multi-Value database market. The MITS product line contains both the established MITS Discover Business Intelligence suite as well as the new report generator, MITS Report. Founded in 1996 and headquartered in Seattle, Washington, the company's technologies are used by thousands of organizations, resellers, and systems integrators worldwide. MITS products are firmly entrenched in a wide range of business areas, including manufacturing, distribution, retail sales and services, education, government, healthcare, and insurance as well as in many other industries.



SpoolerPlus For Unidata

Finally, after all these years, there is a generic PICK spooler for UniData.

No more converting all the STARTPTR and SP-ASSIGN commands in your generic PICK software to SETPTR commands. No need to add PRINT ON statements to every PRINT line in every print program.

SpoolerPlus allows you to run your generic PICK software on UniData using the generic PICK printer commands that are embedded in your software or

that you use at TCL — e.g. :STARTSPOOLER, STARTPTR, SP-ASSIGN, SP-EDIT, LISTPTR, LISTPEQS, etc. — and they behave just like they do on your generic PICK database.

SpoolerPlus will hold, print, hold-and-print and spool on-hold print jobs from designated print queues to designated printers when those jobs are generated and/or when those printers are available.

And (adding a feature to the generic PICK spooler) you can configure each printer to the print characteristics of the assigned print queue automatically when the printer is assigned to that queue, using printer drivers that we supply and/or that you may write — which is the “Plus” in “SpoolerPlus”.

Migrating generic PICK software to UniData has never been easier.

You no longer have to overcome the FILE-SAVE and ACCOUNT-SAVE tape format incompatibilities between UniData and most generic PICK databases when transferring your software. Our FileCaddy programs, TAPE.DUMP and TAPE.LOAD, quickly and reliably transfer all the software and data files in each account from most generic PICK databases to UniData in just two steps, automatically creating new files on the fly if they don't already exist.

The last barriers to moving your software to UniData have been removed.

There's never been a better time to migrate your software to UniData. Using our migration

package, you can migrate your software yourself, we can migrate it for you or we can work with you as a member of your migration team.

Our migration package includes both SpoolerPlus and FileCaddy.

Special pricing for UniData is available as part of our migration package and/or migration services.

To find out how easy it is to migrate your generic PICK software to a database that's adding tomorrow's features today, and that's supported by IBM, contact Dave now at Sysmark Information Systems, Inc., 800-SYSMARK (800-797-6275), davet@sysmarkinfo.com or www.sysmarkinfo.com.

Sysmark Information Systems, Inc. is an authorized IBM Business Partner for the U2 product line, including UniVerse and UniData.



Break Out and Extend your MultiValue Application

Kore Technologies can help you extend and modernize your application with our powerful integration tool suite and our flexible .NET Web solutions. Breathe new life into your MultiValue system...

Integration

- Microsoft SQL Server Databases
- Business-to-Business eCommerce
- Best-in-Class Enterprise Applications
- Business Intelligence and Web Reporting

Web Solutions

- eCommerce Storefront and Portal
- Sales Force Automation
- Customer Service and Help Desk
- Project Management and Collaboration

Kore Technologies is a leading provider of enterprise integration products and services for MultiValue systems. Our solutions, tools, and years of industry experience enable you to take advantage of new technology with your application, extending its life and maximizing your return on investment.

To learn more, call 866-763-KORE or visit www.koretech.com today!

solutions that work.
people who care.



Kore Technologies, LLC
9323 Chesapeake Dr., Suite D
San Diego, CA 92123
866-763-5673
www.koretech.com

A Backup Strategy in the UniData Environment

BY JON A. KRISTOFFERSON

O V E R V I E W

I am a System Administrator of a 24/7 computer operation.

Our largest problem to date has been creating a solid backup of our UniData database while users are working on the system.

Recent enhancements in UniData and AIX have provided the means of creating a solid backup while users are working with the UniData database.

Solid backup

By solid backup I mean that the backup will be taken with users on the system without creating corrupted images of UniData files on the backup medium. In the past, when creating a backup of our UniData Database, we would do so with users working on the system. Our environment demands that we provide 24 hour a day, 7 day a week access to the computer resources for our users.

The problem with backing up the UniData database with users updating the system is that for some large files the backup grabs the start of a file which contains the header information. By the time the end of the file is on the backup medium, the body of the file has been modified without the updates to the header record of the file. When it comes time to restore this file, UniData will see this as a corrupted file!

Everybody knows, or should know, that for a solid backup you need to make the database quiescent. This can be done by logging out all of the users and shutting down the UniData application. This is good for the backup, but bad for the 24/7 mandate of our operation.

Another aspect about making a solid backup is to define a solid backup as a backup without file corruption. This does not mean the same thing as a backup with referential integrity. Referential integrity is better maintained by using transaction logging, and is outside the scope of this article.

Enhancements Required

The enhancement to UniData required to make a solid backup is the ability to pause and resume the UniData database. This can be done with the "dbpause" and "dbresume" commands. The dbpause command suspends writes to the database files and flushes any modified buffers out to the operating system. After the dbpause command completes, the file system is in a consistent state and can be backed up providing a solid backup.

The dbpause command eliminates the need to log off the users by providing a state of the machines operation that mimics all of the users being logged off. The problem with this is again the 24/7 availability of our environment. The time it takes to backup our data sometimes takes several hours. We cannot deny write access to the UniData database for hours at a time. A method is needed to minimize the downtime for the backup process.

The second enhancement to help us with the backup solution is the filesystem snapshot capability. A snapshot is a way of creating a "picture" of the filesystem at a specific point in time. The snapshot is maintained by AIX. The snapshot is a way of creating a filesystem to backup much like the unused portion of a broken mirror set.

So with these two enhancements we can create a solid backup with minimal impact on the users. How do we do this? By per-

forming the following commands in the proper sequence, we can create a solid backup of our UniData database. The sequence is as follows.

1. Dbpause the UniData Database.
2. Create a snapshot of the filesystem containing the UniData database files.
3. Dbresume the UniData Database.
4. Mount the snapshot filesystem.
5. Backup the snapshot filesystem.
6. Unmount the snapshot filesystem.
7. Discard the snapshot.

Practical Considerations

The real issue here is the time it takes to perform steps 1-3. This time is the actual time the users are unable to write to the UniData database files. In our environment we are using UniData 6.1, AIX 5.3 on IBM RS/6000 model 510 systems. The time we experience to perform steps 1-3 is between 1 and 2 seconds. This time span is so quick that the users have never even noticed the brief pause.

We use cron to run a korn shell script to perform our backups every evening. We also use rsync for backing up our data. Rsync is used to make a mirror copy of the snapshot filesystem onto another computer system. A dedicated Kuro Box/HG (Linux) system is used to keep the rsync backups.

The Kuro Box/HG system cost us about \$400 to assemble. In our environment the 250 GB backup capacity of the Kuro Box/HG was enough to store a daily backup, a backup for each day of the week (7), a monthly backup and a yearly backup. The cost of nine sets of tapes required to do these backups was in excess of \$400.

What about Mirroring?

About now you may be asking yourself why use snapshots instead mirroring. My answer to that is snapshots have almost no overhead when you are done with the backup. Upon completion of the backup the snapshot filesystem is unmounted and the



Visage

The Most Modern and Cost Effective Tool Available in the MultiValue Market

Choose the MultiValue Development Environment That Has All The Tools!

- Rapid Application Design
- Business Intelligence
- Report Generator
- email Processor
- Forms Generator
- Fax Gateway

Take a FREE
Visage Test
Drive Today at
www.stamina.com.au

process is completed. If we were using a set of mirrored disks, we would still need to synchronize the mirrors. This can be both time consuming and a very large load on the system to complete. Snapshots do not have this overhead.

An ironic thing about snapshots and mirroring is that the filesystem we snapshot is a mirrored pair of disks! We use mirroring to provide for High Availability and Snapshots

to maximize the uptime for the users while minimizing the overhead of the backup.

How do Snapshots work?

Snapshots are maintained by the operating system. The filesystem used as the source of the snapshot is called the "snapped" filesystem. The "snapshot" filesystem is created on

Continues on page 34

Backup Strategy

Continued from page 33

a “backing store.” A simple explanation of how snapshots work is as follows.

The operating system will first pause writes to the disk, allows those write operations already in progress to finish and flushes all I/O buffers out to the disk. This is similar to the “sync” command.

The next step is to create a snapshot file on the backing store. The backing store is a special area where the operating system keeps the snapshot filesystem data. The snapshot filesystem keeps a copy of the snapped file system metadata. The metadata is the inode and directory information of the file system being snapped.

From the time the snapshot filesystem is completed, the operating system will keep the snapshot filesystem updated any time the snapped filesystem disk is modified. When a disk block on the snapped filesystem is modified, only the original disk block is copied from the snapshot filesystem to the snapped filesystem backing store.

The snapshot filesystem only needs to maintain the original disk block from the snapped filesystem. If a disk block is modified repeatedly, only the original disk block is required to be maintained. This means that only a small percentage of the original snapped filesystem needs to be maintained on the snapshot file system. Typically this is about 10-20% of the original filesystem. This of course depends on the amount of updating that takes place while the snapshot is maintained.

Why use Rsync?

Rsync was originally designed to keep directory structures in sync over slow communication lines. The method rsync uses can be simple described as comparing a master directory structure with a slave directory structure and performing those changes on the slave directory structure to keep both identical. In practice it is a bit more complicated but not much. By using a high speed communications link, like a Gigabit link, rsync is very efficient.

Using rsync, our backup time is about 30-40 minutes, depending on the amount of changes required to make the directory structure on the Kuro Box/HG identical to the snapshot filesystem with the UniData database on it. Our old backup method would take 4-5 hours to our tape drive.

To maximize the rsync performance, the rsync is performed to a “daily” directory every day on our Kuro Box/HG. Then, later the same day, the Kuro Box/HG runs another rsync process to update the daily directory into a directory named for the day of the week. By using the same daily directory every day, we only need to rsync changes that have occurred to the directory structure within the last 24 hours.

The Kuro Box/HG is dedicated to the backup process and can afford the time and overhead of running rsync on the directories that are 7 (daily), 31 (monthly) or 365 (yearly) days old. The cron process performs the rsync operations on the Kuro Box/HG. Special cron jobs are also added for monthly and yearly backups.

Restoring

Another good reason to use rsync is in the restoring of files from our backup. We can simply use a network copy command to move the files from our Kuro Box/HG to our production RS/6000 system. By using links on the Kuro Box/HG system the directory pathnames can appear identical, at least for the daily backups. To retrieve a file from backup you would need a command like...

```
# scp -r root@rmsbkup:/usr/ud/PM100/DAILY .
```

The backups are online and always available.

Conclusion

The latest enhancements to UniData and AIX allow for a solid backup with minimal impact to the users. We have moved from a 4-5 hour backup, where there was a possibility of backing up corrupted files, to a 30 minute backup where the file system and UniData database were quiescent with only a 1-2 second outage to the users.

The backups are on-line and available all of the time. It is a simple network copy command to retrieve any file from a daily backup.

The rsync command is open source and is available on the AIX Linux Toolkit CD-ROM and is free. The Kuro Box/HG is a low cost PowerPC based system. We purchased ours and added Gentoo Linux and a 250 GB disk drive. Our total out of pocket costs were \$400 and 8 hours of time to configure the Kuro Box/HG and Gentoo Linux. [15](http://www.jonkristofferson.com)

Example AIX scripts to create this kind of backup can be found at the Download link at <http://www.jonkristofferson.com>



JON A KRISTOFFERSON

has 30 years experience in the computer business. Currently

he is working in an Enhanced 911 Call Center outside of Chicago, IL. He also works as the Chief Technical Officer for Sumo Computer (<http://www.sumocomputer.com>) which specializes in selling configured Kuro Box/HG system with Gentoo Linux.

Mr. Kristofferson enjoys giving talks on

MultiValue Database issues, Linux, and Open

Source software as well as tech-

nology trends. His background

includes: Electronic technician; Electrical Engineer; Systems Analyst; Principal Consultant and Executive Manger.

In his spare time he enjoys cooking, astronomy, and spending time with his 7 grandchildren.

Overview of mv.NET Version 3

mv.NET consists of three major components: Core Objects, Binding Objects and Adapter Objects.

Core Objects not only provides the underlying framework upon which all other mv.NET components are based, it also provides, in its own right, a wealth of developer centric features enabling the rapid creation of feature rich, high-performance applications using the powerful tools provided by Microsoft's .NET environment. Core Objects has strong integration with Microsoft's Visual Studio, allowing the developer to carry out virtually all aspects of application creation from within the Visual Studio environment. Each class within the Core Objects class library contains a range of interface members — methods, properties and events. It is by the use of these members that the developer is able to carry out a wide range of MultiValue database-related tasks within the .NET environment. Many of the features provided by these classes will be familiar to the MultiValue developer; in fact, wherever possible, Core Objects' classes have been designed to mirror the traditional functionality provided by DataBASIC and other core MultiValue components. However, by virtue of the fact that the client/server architecture differs in a number of fundamental ways from the legacy green-screen application paradigm, mv.NET's classes provide features that will be new to the MultiValue developer.

Binding Objects complements the Core Object component by providing state-of-the-art databinding technology, allowing

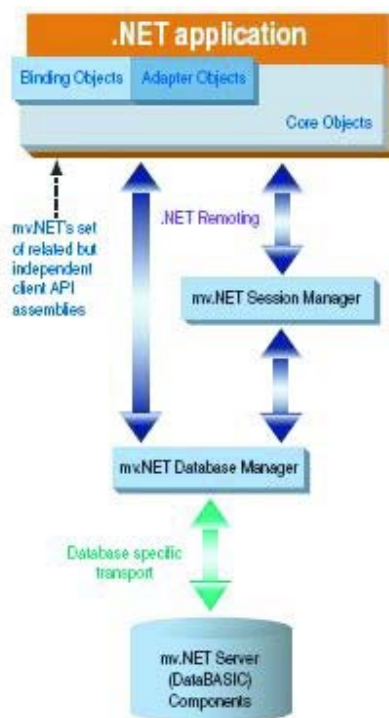
rapid application development from within the .NET environment. All aspects of Binding Objects are fully MultiValue data structure aware, right down to subvalue level. From the simplest file maintenance form to the most complex multi-database, multi-file transaction screen, Binding Objects provides the sophistication capable of meeting the most demanding needs of today's business application developer. More so, it does this in an easy to use, flexible and extendible manner, providing the developer with a databinding technology which provides a richer, more intuitive methodology of mapping the unique capabilities of the MultiValue data model to the user interface; thus allowing more sophisticated use of the MultiValue data model and a greater range of business application development features. Binding Objects components can be used outside of the Visual Studio environment, but really come into their own when used in conjunction with mv.NET's extensive Visual Studio add-in technology. Dockable toolbars providing databinding summary information along with sophisticated auto positioning of automatically created bound controls enable the developer to save valuable time in their task of creating state-of-the-art applications.

Adapter Objects is a full ADO.NET managed data provider implementation and provides a full read/write interface for Mul-

tiValue databases. MultiValue data is dynamically normalized into ADO.NET datasets making this the way forward for true SQL-like application development and reporting, for example, using products like Crystal Reports. To further extend its ease of use, Adapter Objects provides optional support for the use of SQL syntax within Select, Update, Insert and Delete command

objects. Adapter Objects also provides exceptional flexibility when handling MultiValue data, offering a range of different strategies to suit a variety of application requirements. Complementing this flexible MultiValue data handling is a powerful multivalue/subvalue filtering capability, allowing filtering of item data right down to subvalue level across multiple column selection criteria. As a true ADO.NET provider, all datasets and datatables

generated via Adapter Objects may be used with the standard .NET components, such as typed dataset generation and databinding. Adapter Objects also utilizes Core Objects' advanced session management functionality, allowing rapid connect and disconnect of database sessions



Continues on page 45

AJAX & Opensight

From Buzzword to understanding

The “buzzword” part

When I was kid growing up in the suburbs, my mom used AJAX[®] branded products to clean the clothes, the bath, the floor, well just about everything! I was more an ACME brand guy, something to do with Wile E. Coyote I guess.

Anyway, fast-forward 30 years and now I’m using AJAX, not the cleaning products, but the latest buzzword to hit cyberspace. AJAX, the buzzword, stands for Asynchronous JavaScript And XML.

So what’s so new about JavaScript and XML? Not much really. They’re the same as they were last week when AJAX was still a cleaning product. The important letter in this acronym is the first “A”, which stands for Asynchronous. This clever little letter is the reason Internet companies in the know, like Google, are using AJAX to build their latest and greatest web based applications.

But you’ve heard all this before. Web Portals were it, a bit and a packet of chips for a while. What about the Push-Pull and dot Net? So why is AJAX any different? Why should you spend your valuable time on this latest development craze? Because this stuff is easy, it actually works, and you can start doing it now with the tools and knowledge you already have!

The “understanding” part

Since the early days of the Internet, developers have searched for a standard and browser independent way to asynchronously load content on an existing web page without requiring a full reload of the page. Microsoft made early attempts in 1996 when they introduced the `iframe` element into Internet Explorer 3.0. Netscape responded when they introduced the now defunct `layer` element into Netscape 4.

Over the last 10 years techniques like Remote Scripting (1998) and custom Java applets have come and gone. The web development community, collaborating through newsgroups and blogs,

refined these techniques until, in 2002, a new standard was created called XMLHttpRequest. This standard is a set of API's that can be used by many web-browser scripting languages to transfer and manipulate XML data to and from a server using HTTP, and thus establishing an independent connection channel between the client and the server.

XMLHttpRequest is the glue that makes AJAX work. It's that important first "A".

AJAX is not a product or a language; it is a methodology for building interactive web applications. It makes web pages feel more responsive by using JavaScript and XML to send and receive small amounts of data between a browser and the server behind the scenes, therefore removing the requirement for the web page to be fully reloaded to view any changes. This helps to make web applications smaller, faster and more user friendly.

AJAX is based on the following open web standards:

- XHTML (eXtensible HyperText Markup Language)
- JavaScript
- XML (eXtensible Markup Language)
- CSS (Cascading Style Sheets)

XHTML is compatible with the W3C standard HTML 4.01 and will eventually replace HTML. It is supported by all new browsers including Internet Explorer, Firefox, and Opera on most operating systems.

XHTML, in combination with CSS, controls the styling of the web page.

JavaScript is used as a client-side scripting language to interact with the HTML DOM (Document Object Model) to dynamically display and interact with the user and the information presented. Other client-side scripting languages can be used (like PHP) as well as server-side scripting like ASP, although a client-side language is required to interact with the user via their browser.

INTERNATIONAL
SPECTRUM
THE BUSINESS COMPUTER MAGAZINE


Have you seen it?
Everyone's talking about it, International Spectrum e-Xtra™, the new e-newspaper.

Look for e-Xtra, on the first and the fifteenth of each month. This e-newspaper has timely industry news, hot topic, classified ads, and things of interest heard on the MultiValue grapevine.

- Let us know what you've heard on the grapevine that you think we should look into.
- Let us know what's new at your company that you'd like other people to know!
- Find out how to add e-Xtra punch to your *International Spectrum Magazine* ad with a complementing ad in *International Spectrum's e-Xtra*.
- Take advantage of placing an inexpensive classified ad in *International Spectrum e-Xtra*.

Get your FREE subscription at www.intl-spectrum.com

XML is a commonly used format for transferring data from and to the server, although any format will work including preformatted HTML or plain text.

AJAX is web server, browser, and operating system independent. Its applications can reach a larger audience and are easier to install and support than desktop applications.

As you can see, the technologies behind AJAX have been around for several years, but now, under the AJAX banner, have become more widely known and adopted. "Google Earth" and "Google GMAIL" are examples of web applications based on the AJAX methodology.

How AJAX works

A traditional web application will submit input (using an HTML form) to a web server. After the web server has processed the data, it will return a completely new web page to the user.

Because the server returns a new web page each time the user submits input, traditional web applications often run slowly and tend to be less user-friendly.

With AJAX, web applications can send and retrieve data without reloading the whole web page. This is done by sending an XMLHttpRequest to the server and by modifying only parts of the web page using JavaScript.

The preferred way to communicate with the server is by sending and receiving data as XML, but other methods can be used.

Using AJAX with MultiValue Databases (e.g. OpenInsight)

MultiValued databases like OpenInsight work easily with XML. This is because XML is delimited data, but instead of using field (attribute) marks and value marks, XML uses tags.

Continues on page 38

OpenInsight, for example, has built in functions to send and receive XML data. This makes it very easy to use OpenInsight as the database in an AJAX application.

Below we have a simple example of using AJAX techniques to receive data from OpenInsight and display it in an HTML page.

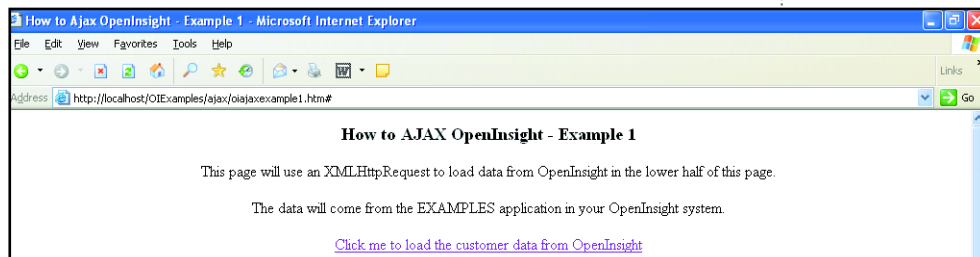
In this example we will be using OECGI.EXE as the means of connecting to OpenInsight from the web server. For more information on setting up and using OECGI.EXE please refer to the OpenInsight help files or the knowledge base section on www.revelation.com.

The example will request the lname, fname, city, state, and zip columns from the customer table, found in the OpenInsight EXAMPLES application, sorting the data by lname, and then display this on the HTML page.

FIGURE 1

```
<html>
<head>
<title>How to Ajax OpenInsight</title>
</head>
<body>
<h3 align="center">How to Ajax OpenInsight</h3>
<p align="center">
This page will use an XMLHttpRequest to load data from OpenInsight in the lower half of this page.<br />
The data will come from the EXAMPLES application in your OpenInsight system.
</p>
<p align="center"><a href="#" onclick="getXMLData()">Click me to load the customer data from
OpenInsight</a></p>
<hr />
<span id="datafromoi"></span>
</body>
</html>
```

FIGURE 2



MultiValued databases like OpenInsight work easily with XML. This is because XML is delimited data, but instead of using field (attribute) marks and value marks, XML uses tags.

FIGURE 3

```
function getXMLData() {
    var url = "http://localhost/cgi-bin/oecgi.exe/inet_xml?cmd=list customers by lname lname
fname company phone fax email"
    loadXMLDoc(url)
}
```

How to AJAX your OpenSight

Step 1. - Create an HTML page.

Create an HTML page that contains the code shown in Figure 1.

This will produce a page that appears in the browser as seen in Figure 2.

Step 2 - Add JavaScript to request the XML data from OpenInsight

We need a few simple JavaScript functions to perform the connection to OpenInsight and retrieve the XML data. This code needs to be added inside the head element of your html page.

JavaScript Function number 1 - getXMLData

In the first function, getXMLData (Figure 3), we set up the URL we need to pass to the web server to retrieve the XML data. In this case, we are using

OECGI.EXE to make the connection to OpenInsight. We call this function from the onload event in the body tag of the HTML page above.

OECGLXEXE makes a call to the BASIC+ function INET_XML. The source code for INET_XML can be found in the SYSPROG account of your OpenInsight system.

We pass a standard RLIST statement to INET_XML via OECGI.EXE, which it uses to select the table, rows, and columns to return as an XML document.

FIGURE 4

```
function loadXMLDoc(url) {
// code for Mozilla, etc.
if (window.XMLHttpRequest)
{
xmlhttp=new XMLHttpRequest()
xmlhttp.onreadystatechange=state_Change
xmlhttp.open("GET",url,true)
xmlhttp.send(null)
}
// code for IE
else if (window.ActiveXObject)
{
xmlhttp=new ActiveXObject("Microsoft.XMLHTTP")
if (xmlhttp)
{
xmlhttp.onreadystatechange=state_ChangeForm
xmlhttp.open("GET",url,true)
xmlhttp.send()
}
}
}
```

FIGURE 5

```
function state_Change() {
// if xmlhttp shows "loaded"
if (xmlhttp.readyState==4)
{
// if "OK"
if (xmlhttp.status==200)
{
//alert("XML data OK")
//Create the xmldoc object and load it
// with the data returned from the server
xmldoc=new ActiveXObject("Microsoft.XMLDOM")
xmldoc.async="false"
xmldoc.loadXML(xmlhttp.responseText)

//Loop through the list of rows returned
var myhtml = ""
var rows = xmldoc.selectNodes("/rows/row")
for (i=0;i<rows.length;i++)
{
var thisrow = rows.item(i)

// Count the number of columns
var cnt=thisrow.childNodes.length
for (j=0;j<cnt;j++) {
// Write the current row to the form
myhtml = thisrow.childNodes(j).text + " "
}
myhtml = myhtml + "<br/>"
}
}
}
else
{
alert("Problem retrieving XML data: " + xmlhttp.statusText)
}
}
}
```



Business and Computer Solutions For Your Real World Challenges

<u>Business Services</u>	<u>Verticals We Serve</u>
<ul style="list-style-type: none"> • Mentoring • Streamlining • Documentation • Training • Process Analysis • Needs Analysis • RFP Writing • Many more... 	<ul style="list-style-type: none"> • Manufacturing • Importing • Exporting • Transportation • Courier • Education • Fund Raising • Retail • Distribution • Fulfillment • Insurance • Not-For-Profit • Health • And Others...
<u>Programming</u>	
<ul style="list-style-type: none"> • Web • 4GL • Object Oriented • Green Screen 	

www.KeyAlly.com
Results@KeyAlly.com
(718) 762-3884 x 1

Key Ally, Inc.
P. O. Box 540957
Queens, NY 11354, USA

**NEEDED
PICK PROGRAMMERS**

We have a number of requests for programmers with experience using PICK, Universe, Unidata, SB+, jBase or Advanced Revelation from client companies nationwide.

**Email your resume to: Chuck Okeson at
Chuck@softwareSearch.com**

800-949-5423 x203

Have a new software product for the MultiValue market?

Combine the power of your AD in **Spectrum magazine** with a marketing campaign to one of our qualified email lists!

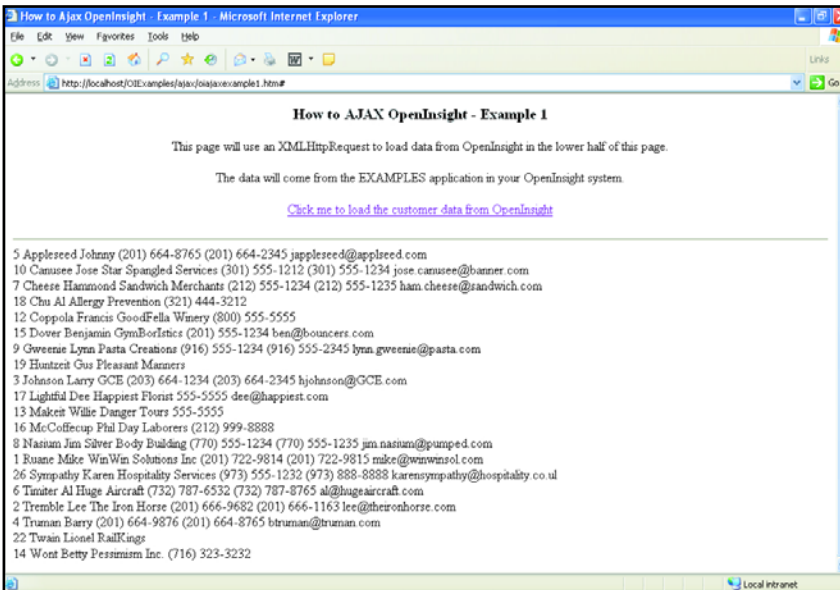
It's more affordable than you think.

Call **720-259-1356**
for a free quote today.

AJAX & Opensight

Continued from page 38

FIGURE 6



XMLHttpRequest makes an asynchronous connection to the server. This means that it sends the request, but does not wait for a response.

For Example:

The following call to OECGI.EXE will start the inet_xml BASIC+ routine. It will take the Rlist command specified by the cmd= parameter and select the customers table by the lname column, returning the lname, fname city state city columns as an XML document.

Eg: `http://localhost/cgi-bin/oecgi.exe/inet_xml?cmd=list customers by lname fname lname city state zip`

The getXMLData function then calls a JavaScript function called loadXMLDoc which will make the connection to Opensight to get the data.

JavaScript Function No 2 – loadXMLDoc

The loadXMLDoc function uses the XMLHttpRequest object to make a connection to the web server to request the data.

Please note that the syntax used to setup the XMLHttpRequest varies between browsers. The example below handles IE and Mozilla browser types.

XMLHttpRequest makes an asynchronous connection to the server. This means that it sends the request, but does not wait for a response. Instead, we have to tell XMLHttpRequest the name of a function to call when the state of our request changes.

This is done by setting the onreadystatechange property of the XMLHttpRequest object.

In the example shown in Figure 4 on page 39, we set this to call a function called status_Change. As the status changes for our XML request, this function will be called. (See step 3 for more information)

Step 3 – Capture the XML Document in the state_Change function

This is where all the magic happens. The XMLHttpRequest state_Change event (shown in Figure 5 on page 39) will call the specified function whenever the state of the XMLHttpRequest changes. The different states are as follows:

- 0 = uninitialized
- 1 = loading
- 2 = loaded
- 3 = interactive
- 4 = complete

Once the state value is 4, the XML data has arrived back at the browser. Any errors are returned in the status and statusText properties.

The stateChanged function then initialises an XMLDocument object to hold the XML data returned from the web server. Once we have an XML document object initialized, we can start displaying the XML data in our HTML page.

In the example above we loop through each row returned and display the fname and lname columns down the

page inside the span tag. The results are shown in Figure 6.

A quick note on the XML Document Object Model (DOM)

The XML document object model provides a text-based means to describe and apply a tree-based structure to information. There are many ways to read this tree structure. XPATH is one way to navigate and extract information in an XML document.

XPATH is a major element in the W3C XML standards as so is fundamental to a lot of advanced XML usage. Please visit www.w3schools.com for more information on XPATH.

Putting it all together

Figure 7 shows the complete html page containing all the elements required to get you up and going with your first AJAX enabled web page. The source code for this page and another two examples are available to download from the knowledge base section on www.revelation.com. Search the site for the word AJAX.

So, again we can strike up another victory for our little MultiValued product. The market has come up with yet another, "all new, improved" buzzword, and we are buzzword compliant with just a little bit of work. Will it always be this easy? Maybe not. But OpenInsight can give you the AJAX functionality required by customers in a quick and easy fashion. [15](#)

Note:

The complete html page containing all the elements required to get you up and going with your first AJAX enabled web page can be found at the Revelation web site. The source code for this page and another two examples are available to download from the knowledge base section on www.revelation.com. Search the site for the word AJAX.

```
<!doctype html public "-//w3c//dtd html 4.0 transitional//en">
<html>
<head>
<title> How to Ajax OpenInsight - Example 1</title>
<meta name="generator" content="openinsight">
<meta name="author" content="David Goddard">
<meta name="keywords" content="way cool">
<meta name="description" content="OpenInsight,AJAX">
<script type="text/javascript">
function getXMLData() {
    var url = "http://localhost/cgi-bin/oecgi.exe/inet_xml?cmd=list customers by lname
lname fname company phone fax email"
    loadXMLDoc(url)
}

function loadXMLDoc(url) {
    // code for Mozilla, etc.
    if (window.XMLHttpRequest)
    {
        xmlhttp=new XMLHttpRequest()
        xmlhttp.onreadystatechange=state_Change
        xmlhttp.open("GET",url,true)
        xmlhttp.send(null)
    }
    // code for IE
    else if (window.ActiveXObject)
    {
        xmlhttp=new ActiveXObject("Microsoft.XMLHTTP")
        if (xmlhttp)
        {
            xmlhttp.onreadystatechange=state_Change
            xmlhttp.open("GET",url,true)
            xmlhttp.send()
        }
    }
}

function state_Change() {
    // if xmlhttp shows "loaded"
    if (xmlhttp.readyState==4)
    {
        // if "OK"
        if (xmlhttp.status==200)
        {
            //alert("XML data OK")
            //Create the xmldoc object and load it
            // with the data returned from the server
            xmldoc=new ActiveXObject("Microsoft.XMLDOM")
            xmldoc.async="false"
            xmldoc.loadXML(xmlhttp.responseText)

            //Loop through the list of rows returned
            // We are using an XPATH statement to select all the rows
            // returned in the XML document.
            var rows = xmldoc.selectNodes("/rows/row")
            var myhtml = ""
            for (i=0;i<rows.length;i++)
            {
                //var thisrow=xmldoc.documentElement.childNodes(i)
                var thisrow = rows.item(i)

                // Count the number of columns in the row
                var cnt=thisrow.childNodes.length
                for (j=0;j<cnt;j++)
                {
                    // Write the current row to the form
                    myhtml = myhtml + thisrow.childNodes(j).text + " "
                }
                myhtml = myhtml + "<br />"
            }
            document.getElementById("datafromoi").innerHTML=myhtml
        }
        else
        {
            alert("Problem retrieving XML data:" + xmlhttp.statusText)
        }
    }
}
</script>
</head>
<body>
```


4 GL Briefing — IBM SB+

B Y C H A R L E S B A R O U C H

We asked IBM for a sit-down to discuss IBM SB+ (System Builder), and they certainly responded. Susie Siegesmund, (Director, U2 Data Services), Jackie Burhans, (U2 Portfolio Manager), Thom Garcia, (U2 Tools Product Manager), and Ehab AbuShmais, (U2 System Builder Advanced Technical Support Engineer) all joined us on a conference call to talk about their 4 GL, its capabilities, and its future.

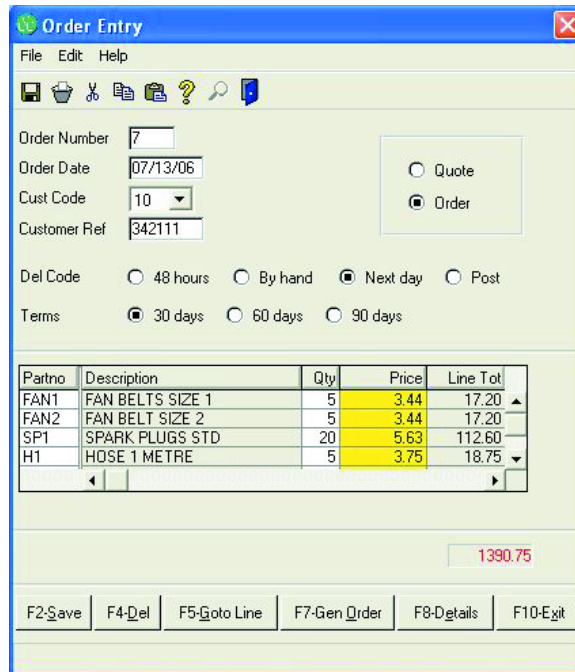


FIGURE 1 - Order Entry screen in IBM SB+

automated regression test suite has made our Quality Assurance much stronger than it used to be. We've updated the look and feel of the products. We feel that IBM SB+ as an underlying development language is a very robust language.

SPECTRUM: I know you have big plans for the future, but what compelling features does it have today?

Jackie: The whole concept of being process-based, the whole concept of reusability, it just makes it possible to build robust business applications very, very quickly.

SPECTRUM: How long has System Builder been out there?

Jackie: About 1995, around that time frame, UniData acquired System Builder, the company. I'd have to check the dates.

Susie: I can work it backwards. We became IBM in July 1, 2001. We became Informix on March 1, 2000. That means we became Ardent Software in February of '98. The product has been with us all that time. And was originally developed by the Australian team that created it.

Moving forward, we've brought things up more tightly to IBM standards. Our

Ehab: The power of it is the ease of use. If you have the requirements, we provide the mechanism for the field definitions, the screen developer, the report writer. Creating applications can take literally minutes.

Susie: While we are looking to keep it up with the demand of the modern world, we feel that the product is already very robust where it stands. For example, the next release will have better multi-byte functionality, since we see more and more partners going into the Orient.

SPECTRUM: If you are new to SB+, what is the learning curve?

Ehab: This is what I'm always trying to highlight, ease of use. You don't even need to learn the command language or learn BASIC. It will prompt you, and we'll generate the BASIC code for you. Very intuitive.

SPECTRUM: What's next for SB+?

Jackie: We just had a very well attended webcast about SB+ 5.4. There were a lot of new features to discuss.

Thom: I'll recap. We have a number of features coming out. One of the more exciting ones is that we will have the IBM U2 Web Development Environment (formerly Redback) interface that will let you convert SB+ processes to web pages. For example, we have a client who moved their SB+ mortgage application process to the web. From an end user perspective, the application is on the web. From a developer perspective, it's an SB+ application. From loan origination to document writing, all of that used to be done in

the background. We've built a platform to support all of that online.

Jackie: We have very active plans for SB+ and the tools around it. We expect it to continue to be a tool that our partners use to develop a competitive edge.

Susie: It's an excellent tool for developing new applications, especially as we continue to add new IBM technologies. Also, this is the last version for the current client. That's because we have a new client for SB+ under development. We are building the new one from the ground up.

Thom: Also, our partners will be able to send out application installers that know how to customize SB+ for the end users requirements. We've added Sarbanes-Oxley audit tracking options, licensing changes, including a sixty day evaluation option. We have removed the dual log in requirement.

SPECTRUM: When will the evaluations become available?

Susie: It will go up on the site after the release of the product. The target for the full 5.4 release is May 31st. So some time in early June for that version. If you want to build something that you are going to resell, you become an IBM OEM partner. We use the OEM (Original Equipment Manufacturer) term, just like it's used in hardware.

SPECTRUM: How do you become a partner?

Susie: There are no fees to become an IBM partner. If they are going to be a support providing partner, you have to have two certified analysts or developers. We use an earned discount model, so as our partners pass certain volume thresholds, they get better pricing.

SPECTRUM: What sort of partners do you already have? Who's using System Builder?

Ehab: Well, without mentioning individual customers, distribution, manu-

Continues on page 44

Manage UNIX Security for uniVerse, UniData, and PI/open the Easy Way

SENTRY

The Integrated Security System

- Simplify administration of users, groups, file permissions, and ACLs
- Protect database commands
- Audit every attempt to use prohibited database commands and objects
- Locate inconsistencies and errors in users, groups, and permissions
- Allow users to change passwords without having access to the TCL or UNIX prompt
- In-person installation and training included in North America

FITZGERALD & LONG

Phone: (303) 755-1102
www.fitzlong.com

4 GL Briefing — IBM SB+

Continued from page 43

facturing, the mortgage company I mentioned, library systems, one of Ticketmaster's competitors...

Susie: Call support centers, we have a benchmark on that — we scaled up to 15,200 users (simulated) against over a billion records without a noticeable slow down in user performance. They turned out the lights on us and we had to go home. I think the answer is that it's applicable wherever you use U2, and you can use U2 everywhere. is

CHARLES BAROUCH is the CTO for Key Ally, Inc. and Mount Olympus Systems, Inc. He is also President of the U2UG and a regular Spectrum Magazine contributor.

Charles Barouch Results@KeyAlly.com - Consulting
Zeus@MtOlympus.us - ETL/Migration/Integration
(718) 762-3884x1

SB+ HISTORY

- 1987** System Builder Corporation founded with System Builder 4GL tool
- 1990** SB+, new generation 4GL tool released
- 1995** Unidata, Inc. acquired SystemBuilder
- 1998** Unidata and Vmark merge to form Ardent
- 2000** Informix acquired Ardent including U2 and SB+
Informix spins off E-Team company and contracts out SB+ development
- 2001** IBM acquired Informix assets including U2 and SB+
- 2003** IBM U2 brings SB+ engineering back in-house to Denver
Numerous point releases with features, fixes and branding improvement
- 2006** Scheduled release of IBM SystemBuilder 5.4 and IBM U2 Web Development Environment for UniData and UniVerse 4.3 with support for SB+ screen migration to the Web.





mv://e-Store

You can have e-Commerce, too.
..Be a part of it today...

eagle rock information systems



info @ eriscorp.com

LI	Model / Part Number	Price*	Quantity	Ext. Price	Instock**
1	PD49VP6C200Q-68794	\$399.00	1	\$399.00	NO
	Part Number: 68794	Update Delete			
	Description: Ruggedized PicoDot: Convergent Laser				
2	PD49VP6LLP W/30'	\$345.00	1	\$345.00	NO
	Part Number: 68528	Update Delete			
	Description: Ruggedized PicoDot: Polarized Retro Laser				
3		\$0.00		\$0.00	YES

Overview of mv.NET Version 3

Continued from page 35

Version 3 Enhancements

Core Objects

Support for additional MultiValue databases

The list of supported MultiValue database platforms has been extended - enabling an even wider range of developers to utilize its powerful feature set.

Support for additional socket-based database connectivity options

To complement the existing telnet and socket-based connectivity options, Version 3 adds extra high-performance connectivity capabilities.

Web Service Based Remote Access Remote connectivity via web services is introduced in version 3 in the form of mv.NET's new 'Gateway' feature. This allows VPN-less connections for remote users to be established via firewall friendly http-based web services.

Windows component/application invoking from DataBASIC

Version 3 adds the ability for DataBASIC code to trigger actions within the Windows environment. This is perfect for situations such as the need to dynamically synchronize server data with, say, Outlook mailbox data, or the ability to launch Windows-based document creation and printing from within green-screen applications.

Extended Visual Studio .NET integrated help

The mv.NET plug-in functionality has been substantially enhanced with VS.NET integrated help to provide greater intellisense and dynamic help support within the product.

Binding Objects

Binding Objects support for Web applications

All of the advantages of mv.NET Binding Objects technology is now extended to the development of browser based appli-

cations with the introduction of full support for WebForms. This brings all the RAD benefits of faster time to market and easily supported code to web application developments as well as for Rich Client.

Support for additional standard .NET controls

Binding Objects now extends the range of standard .NET controls supported by the Binding Object databinding framework, for example the new DataGridView control in VS2005.

A multi-column, fetch-on-demand combo box

An absolute must for the database-centric application developer, the new mvComboBox control supplied as part of the Binding Objects suite, provides sophisticated multi-columned drop-down lists with fetch-on-demand population of drop down content. This functionality really comes in to its own when used in conjunction with the new DataGridView control in VS2005.

Extensible architecture for 3rd party control integration into Binding Objects

Binding Objects now offers the ability for developers to integrate any 3rd party control into the powerful Binding Objects databinding framework.


Adapter Objects

Full SQL-style query syntax support within Adapter Objects

Version 3 provides optional SQL Select, Update, Insert and Delete command syntax support for developers using the Adapter Objects component set; allowing a broader range of skill sets to utilize the product. Ideal for non MultiValue aware programmers to quickly start work using the full capabilities of ADO.NET functionality.

Enhanced Multivalue "explode" Filtering

The CommandText property in Version 3 has been extended to support the "MVFILTER" keyword. This enables extended support for multivalue/subvalue filtering when data is being retrieved in (exploded) 1st normal form. IS



Continued from page 46

thing; but we don't have the time or luxury to rewrite it."

Admit it. You've had this opinion for years. You've complained about it. You've tried to get Management to admit it and let you rewrite it. Some of us have even perpetuated the problem by writing clandestine, undocumented "new" libraries (see above, ahem). You ended up with a fragile, undocumented application that breaks every time you touch it. You spend nights and weekends fixing data corrupted by the latest production "release" because no one in the organization understands proper testing methodology or will allocate the time for it if they do. So you are no longer going to suffer alone. Now the lack of disciplined coding, ignorance of testing, and disdain for documentation ("use the Source, Luke") chaos can be felt by the rest of the World.

Be prepared for the World to bite you on the posterior. I am not a lawyer (my parents were married). But my layman's prediction is that if this SOA thing takes off, within five years we will be reading about lawsuits where a Service presented to the World at large delivered wrong data or performed the wrong action at the wrong time. And you know something? I'd have to side with the complaining party. If you publish a Service that "contracts" to deliver certain data or perform a particular function, you had better get it right. If you don't, your successor will.

This could get ugly. IS

Agree? Disagree? Want to tell Cliff that

he's got his head screwed on back-

wards? Give us your feedback. E-mail

to editor@intl-spectrum.com.

**Subscribe on-line at
www.intl-spectrum.com
or
fax this form to (603) 250-0664
and keep Spectrum magazines
coming FREE for 1 year!!**

_____ NAME		_____ TITLE	
_____ COMPANY			
_____ ADDRESS			
_____ CITY	_____ STATE	_____ ZIP	
_____ PHONE		_____ FAX	
_____ E-MAIL		_____ WEB SITE	
_____ SIGNATURE			_____ DATE

(All questions must be answered. Incomplete forms will not be processed. Complimentary subscriptions are limited to U.S. addresses.)

1. What is your job function/title?

- | | |
|---|---|
| <input type="checkbox"/> Principal/Owner | <input type="checkbox"/> Sales/Marketing |
| <input type="checkbox"/> President/GM/CEO | <input type="checkbox"/> Programmer/Analyst |
| <input type="checkbox"/> MIS/DP Manager | <input type="checkbox"/> Purchasing |
| <input type="checkbox"/> Controller/Financial | <input type="checkbox"/> Consultant |
| <input type="checkbox"/> VP/Department Head | <input type="checkbox"/> Other _____ |

2. Is your company a (check one):

- | | | |
|---|---|---|
| <input type="checkbox"/> Computer System Supplier | <input type="checkbox"/> Dealer/OEM/VAR | <input type="checkbox"/> Software House |
| <input type="checkbox"/> Consultant | <input type="checkbox"/> End User | <input type="checkbox"/> Other _____ |

3. What MultiValue Databases does your company use? (check all that apply)

- | | | | |
|--------------------------------|--|-----------------------------------|--------------------------------------|
| <input type="checkbox"/> D3 | <input type="checkbox"/> Native MultiValue | <input type="checkbox"/> Reality | <input type="checkbox"/> Other _____ |
| <input type="checkbox"/> jBASE | <input type="checkbox"/> uniData | <input type="checkbox"/> UniVerse | <input type="checkbox"/> uniVision |

5. What major business/industry most clearly describes your company?

- | | | | |
|--|------------------------------------|---|---------------------------------|
| <input type="checkbox"/> Accounting | <input type="checkbox"/> Medical | <input type="checkbox"/> Direct Marketing | <input type="checkbox"/> Legal |
| <input type="checkbox"/> Banking/Finance | <input type="checkbox"/> Dental | <input type="checkbox"/> Construction | <input type="checkbox"/> Retail |
| <input type="checkbox"/> Education | <input type="checkbox"/> Insurance | <input type="checkbox"/> Other _____ | |

6. What are your firm's approximate gross annual sales?

- | | |
|---|--|
| <input type="checkbox"/> Under \$500,000 | <input type="checkbox"/> \$500,000 - \$1 million |
| <input type="checkbox"/> Over \$1 million - \$5 million | <input type="checkbox"/> Over \$5 million - \$10 million |
| <input type="checkbox"/> Over \$10 million - \$25 million | <input type="checkbox"/> Over \$25 million - \$100 million |
| <input type="checkbox"/> Over \$100 million - \$500 million | <input type="checkbox"/> Over \$500 million |

IS 3/06



Continued from page 47

So I admit, I am very skeptical about claims of reusability made for SOA. After all, there may be a service that performs the operation I need, but it's undocumented, and my services always run better than your services anyway. Still, we've been chasing this reusability chimera for several decades now. What will another attempt at it hurt? This time things might actually get documented so people know what is there and how it works. Programmers might just put aside their egos and concentrate on solving business problems rather than writing code that re-implements existing functionality and adds a bunch of really cool stuff that the user community didn't ask for, doesn't want, and won't use. I mean, it could happen.

No. Something else has been bothering me, and I think I finally figured out what it is.

When I read articles about SOA, I see "security concerns" mentioned. A sentence or two later, I see "outside the organization's firewall." Thinking back to the discussions about directories and service metadata and the like, it dawns on me.

"Oh my stars and toggle switches! They intend to publish these services to the World!" (I don't really say, "Oh my stars and toggle switches!" But Spectrum management would never let me print what I really do say. Just as well...) In this case, of course, The World is defined as "outside service consumers." That could be the World at large or just the World of your B2B partners. This idea sends chills up my spine. Why? Follow me on this.

A lot of MultiValue shops are running what they euphemistically call "legacy code." This is a code phrase for, "We have no clue how this stuff works, and it breaks every time we change some-

Continues on page 45

Reaping What You SOA



BY CLIFTON OLIVER

There is something about SOA that bothers me.

I started hearing about Service Oriented Architecture (SOA, pronounced “so-ah” by most English speaking people and Americans) about a year ago. I suppose the first thing that caused some twitching of the needle on my bogometer were the breathless assurances that this was The Thing that was going to revolutionize Information Industry.

For those not familiar with the bogometer, it is a device that measures bogon rates, usually in bogons per second. Bogons are an elementary particle of marketing output. The higher the bogons per second, the higher the bogosity factor of any given industry “trend.” A rating greater than 95 qualifies a statement as “totally bogus.”

After thirty-seven years in the industry, I admit my bogometer is a bit over sensitive. It’s just that it has had its needle full pegged too many times. I was writing assembly language (1401 Autocoder, not that fancy-dancy 360 stuff) when we were assured that high level languages like Fortran and Algol would eliminate coding errors (bugs) by raising us out of the quagmire of machine level details.

A few years later, we were told that the days of the professional programmer were numbered, as COBOL would allow MBAs to write their own pro-

grams in near-English. The fact that most MBAs cannot write in English to begin with went unnoticed.

In the 70s, the demise of the programmers was again predicted as 4GLs would let non-programmers describe their business and automatically generate an application. Yes, they really thought that those people who still cannot tell you what they want or how their business really works would be able to describe it to a computer.

In the 90s, the demise of the programmer was (yawn) yet again predicted since GUI driven database tools like FileMaker and Access would allow business people who couldn’t figure out how to use a 4GL and couldn’t describe their business to an analyst would simply slap together a screen and some reports and have an application ready to use.

Oh, this has worked out well.

Looking at the predictions for SOA, however, I am not finding any such claims. On the contrary, most of the emphasis seems to be on how information management professionals (programmers, to use the archaic term) can deliver services to the business better and faster. So claims that it is going to fundamentally change the industry population must not be what set off the bogon alert. Ah, hah! It must have been the “R” word—reusability.

Maybe this really is the thing that finally makes reusability a reality, but I’m not holding my breath. Subroutine

libraries in high level languages were supposed to dramatically cut the amount of time it took to get things into production and reduce the time it took to test. What you actually saw, however, is that aside from some mathematics, statistics, graphics, and operating system libraries, programmers kept coding the same application things over and over. There may be a library in the shop that performs the functions necessary for this particular business function, but why on earth would someone want to use it? First, it’s undocumented. Most programmers believe they can write code faster than they can understand what someone else has written. Unfortunately, considering the state of most code I’ve seen, they are right. Second, we all know that your library won’t be as good as the one I’m going to write, right?

Then came object-oriented programming. Classes were the new thing, and one of the Big Deals about OOP was that the classes are reusable. They were going to dramatically cut the amount of time it took to get things into production and reduce the time it took to test. Again, there are a number of classes that really do get reused—windowing, sockets, database access, and the like. Again, however, when it comes to business rules and procedures, we see classes being written over and over to do the same business function. There may be a class that does what you want it to do, but first, it’s undocumented. You get the picture.

Continues on page 46

If two of your top customers stopped buying from you... how would you know?

MITS Advanced Reporting and Business Intelligence.

More MultiValue* customers have
turned to MITS for Business
Intelligence than all other products
combined. Why haven't you?

*UniVerse™, UniData™, D3™ and others



mits

Advanced Reporting
and Business Intelligence

Visit MITS today at www.mits.com or contact us at 888 700 6487

© 2006 Management Information Tools, Inc. All rights reserved. MITS, MITS Discover and MITS Report are trademarks or registered trademarks of Management Information Tools, Inc. in the United States and/or other countries. All other company and/or product names may be trademarks or registered trademarks of their respective owners.