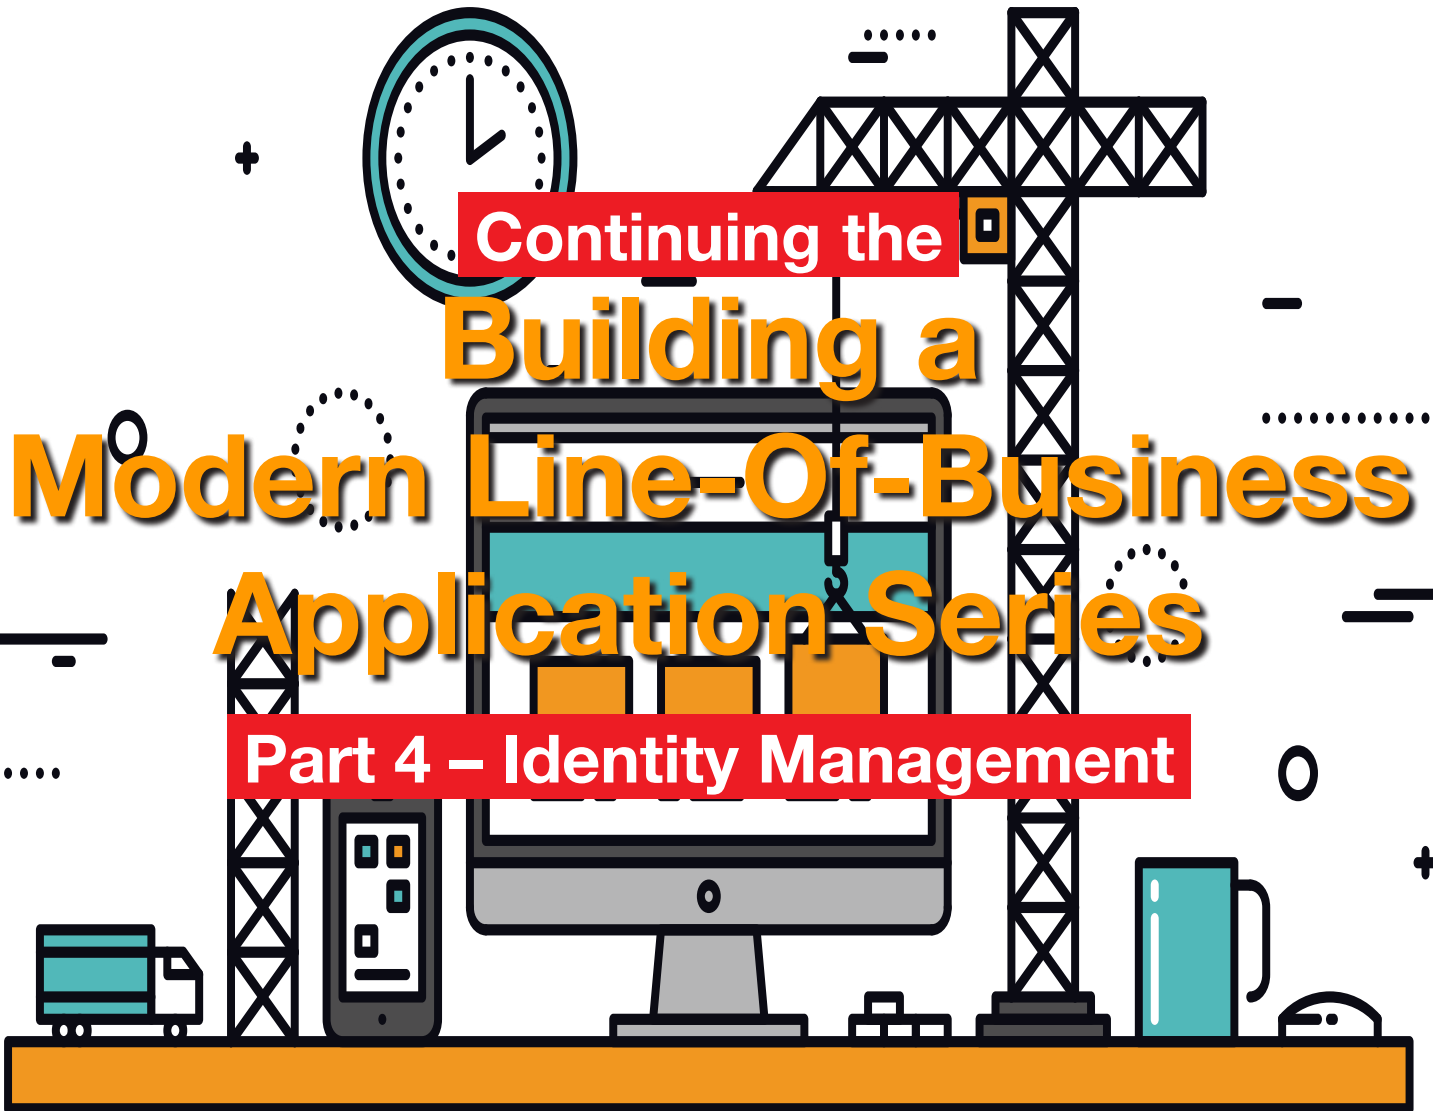


# INTERNATIONAL Spectrum

THE MULTIVALUE  TECHNOLOGY MAGAZINE | JULY/AUGUST 2016



Continuing the

## Building a Modern Line-Of-Business Application Series

Part 4 – Identity Management

### Also In This Issue:

- Save Your Business From Flawgorithms
- QM With a Side of SQL
- Unstructured Data: Do We Have a Home Field Advantage?



INTERNATIONAL

*Spectrum*

36<sup>TH</sup> ANNUAL  
MULTIVALUE CONFERENCE

LITCHFIELD PARK  
ARIZONA

The Premier MultiValue User Event

MAY 1 – 4, 2017

THE WIGWAM

LITCHFIELD PARK, ARIZONA


[www.intl-spectrum.com/conference](http://www.intl-spectrum.com/conference)

The **International Spectrum MultiValue Conference** is the only place that provides diverse solutions and support for integrating and developing your Enterprise applications into the world of business and technology.

This annual event is the place to find the **knowledge, people, resources, and technologies** for enhancing your MultiValue (D3, UniVerse, UniData, Reality, jBase, Caché, QM, and OpenInsight) software applications.

Save the Date!





## 6 Building A Modern Line-Of-Business Application Part 4: Identity Management

This is part 4 in a series on exploring what it takes to build a Line-Of-Business application from scratch using the tools and features found in modern software technologies. In this article we explore identity management. What does it take to provide security, access management, and business discipline to the users of your software? Read Part 4 and find out. **BY NATHAN RECTOR**

## FEATURES | JULY/AUGUST 2016

10 **Business Tech: Flawgorithms** I'm on Facebook, talking to Ross Morrissey, and we end up in a conversation about flawed assumptions in formulas. It started me thinking. **BY CHARLES BAROUCH**

16 **QM and VFS** Sometimes MultiValue can be the complete solution. Sometimes we just use pieces. The database can be used without using mvBASIC or Proc. Or, you can code the middle in MultiValue and let SQL hold the data. Here's a success story with QM can playing that middle role. **BY BRIAN LEACH**

22 **Unstructured Data** In the world of data storage and retrieval, life used to be simple. You had elements of data. Data had various types. Data elements went in columns in a table. And those of us in MultiValue had it even easier because we weren't even bound by those rules. But now we have to deal with unstructured data — email, social media posts, videos, etc. This article explores some of the issues that will cause even a MultiValue Developer to react with, "Oh, my..." **BY NATHAN RECTOR**

## DEPARTMENTS

**From the Inside** page 4

**From the Press Room** page 12

**Master Class in MultiValue** page 21

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

# From the Inside

It is the time of year that I start planning for the next Spectrum conference. If you haven't yet seen the ad for the 2017 Spectrum, it will be May 1st-4th at The Wigwam in Litchfield Park, (Phoenix), Arizona. This is the same venue as the 2016 conference.

It seems like the conference is pretty far away, but it really isn't. I have to start planning everyone's sessions so you have something to show to your company. I would like to know what you, as a MultiValue user, would like to see. What topics will sell your bosses on sending you to Spectrum?

The demands on your IT department are getting more specific, which defines the types of examples and business solutions that you need to know more about. I remember having to justify my attendance, so I watch the trends that are coming. The conference works because we update our content with the times.

We have MultiValue Experts as speakers who give us great content, but it is up to me to guide them. It is up to you to guide me. I would like to hear from the MultiValue Users and Developers. I want to make sure you get the materials, education, and experiences that are needed to maximize your ability to create efficient enterprise applications.

There are a lot of different trends floating around right now. If you are unsure what you would like to see, here are a few that I've noticed:

## Virtualization/Cloud Computing

Everyone is talking Cloud Computing and Virtualization, and for

pretty good reasons. We have presented topics on Virtualization and Cloud Computing at the conference before, but what are you interested in hearing next? Do you need to know how to integrate MultiValue databases into a cloud API, like Azure or Google App Engine? Or are you looking run a MultiValue server in a VM(Virtual Machine)?

## Tablets and Mobile Apps

Mobile Line-of-Business applications are a must. Many companies are not taking full advantage of what mobile apps can bring or how they should be used in the enterprise. How can we help you help them?

## Windows 10 and Universal Apps

Working with Windows 10 development can be a pain, especially if you haven't been part of the beta program. Would you like to see general sessions on how to use Windows 10? On how Universal Apps are created? These would probably not be MultiValue specific, but we need to know about other software to make our jobs possible, so I'd be happy to include them.

## Office 365 REST Services

Looking to integrate Office 365 more closely with your Line-of-Business application? Documents, calendars, e-mail, spreadsheets... there are versions that are free and versions for a fee. All can be integrated with your MultiValue data.

## Integration - Email, CRM, EMR, LDAP, Active Directory

Do you have other applications within your environment that you need to integrate with? We have

NATHAN RECTOR  
*President*

CHARLES BAROUCH  
*Editor*

SYDNEY BAROUCH  
*Editor*

TRACEY RECTOR  
*Layout*



Learn more about the MultiValue Symbol and see what MultiValue Technologies and MultiValue Communities exist to help you support and manage your business and systems. To find out more visit

<http://www.intl-spectrum.com>

---

**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., 3691 E. 102<sup>nd</sup> Ct., Thornton, CO 80229; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: [request@intl-spectrum.com](mailto:request@intl-spectrum.com). Copyright 2016 International Spectrum, Inc. All rights reserved. Reproduction in whole or in part, without written permission, is prohibited.

PRINTED IN USA

## 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 and diagrams are welcome. Although there is no guarantee a submitted article will be published, every article will be considered. Please send your press releases, articles, and queries to: [editor@intl-spectrum.com](mailto:editor@intl-spectrum.com). *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.

**can**  
**MultiValue ~~can't~~ do that.**

**Change your thinking.**

**Modern MultiValue, LLC**

*Modern Solutions for Modern Applications*

info@ModernMultiValue.com

www.ModernMultiValue.com

+1.225.341.1778

been an island unto ourselves for so long, we sometimes forget that our business has information that other systems can use, or is necessary for business functions. What integration options do you need?

#### External Programming Languages

What languages are your business requiring you to work with? .NET, Java, Python, Javascript, HTML, JSON... there are many others. We can provide topics on all of these, both stand-alone and in connection with your databases.

These are just a few examples. What do you find most interesting? What are the trends and topics that you would like to have covered in sessions at Spectrum? I would like to hear from you.



**NATHAN RECTOR**


President

International Spectrum

nathan@intl-spectrum.com

**mv**  
**QB**

## QuickBooks API for the MultiValue Database



- **Read/Write Directly to Quickbooks Databases**  
Customer, Vendor, Invoices, Purchase Orders, Chart of Accounts
- **mvQB API is Designed for the MultiValue Program to Use**  
All routines are simple BASIC calls designed for the developer. No special user interfaces required.
- **No Need to Learn the Internals of QuickBooks**
- **QuickBooks Pro/Premier/Enterprise**

**NATEC**  
**Systems**



Providing Solutions to your MultiValue Questions

**Phone:** 303.465.9616

**E-mail:** mvqb@natecsystems.com

**Website:** www.natecsystems.com

# BUILDING A MODERN LINE-OF-BUSINESS APPLICATION



## Part 4 – Identity Management

### Identity Crisis

Identity Management is often known as User Management. Security Management is often known as Access Management. A lack of common terminology is just one of the issues which makes this topic complex.

What is Identity Management? If you have ever built, or used, a multi-user, multi-interfaced software before, you've encountered it. Identity Management is used to provide security, access management, and business discipline to the users of your software.

Unfortunately, it is not always well thought out. Many developers rely on the built-in features of the database or O/S, and then find themselves boxed in and unable to change to something better or more flexible.

It is the key and firewall to your LOB (Line-of-Business) application. A good Identity Management system allows your application to avoid running into

*What about the IoT (Internet of Things)?*

*If you have a sensor transmitting data to your LOB, you probably don't want it interfacing with your web, GUI, or green screen.*

restrictions that may require difficult workarounds.

### Security

Identity Management is commonly used for security. If your software merely checks for the username and password, you are already missing opportunities. For example, you could be making sure that the user is permitted to access the LOB software from a specific interface or device. Controlling access is no longer a yes-or-no process.

The security aspect of Identity Management must also include validating

security tokens using them to expire old passwords, enforce user idle timers, and authenticate logging.

### Access Management

Access Management is the more sophisticated side of security. It is all about what the user can do once they have been allowed inside. Many developers will equate this to “menu access” or in modern terms “role management.” That's just a part of the story. Access Management should be designed to provide administrators complete control over what a user does, when it can be done, for how long, how much, and/or where it can be done.

One way to accomplish this is by establishing roles or groups. Roles provide the administrator the flexibility to keep things simple for most users while still allowing for complex adjustments where needed. This is how Microsoft ActiveDirectory works, even if it isn't the best designed example.

## Business Discipline and Procedures

This may seem like an odd idea to single out, as it just seems like an obvious extension of Access Management. In its purest form that is correct. A developer would use a role to define whether or not the user has been granted access to a specific procedure or process. But your system can also define if a user can change information. Or perhaps, if they are required to answer additional questions when dealing with data. You could also use this system to define what and how specific users or interface will enter information.

The simplest example is assuming certain data (customer id, terms, credit limits, or available products) when working with a mobile device, while explicitly asking those questions from the user when working at a desktop.

There are certainly people proficient in texting, but the general rule is that less typing is preferable when dealing with mobile.

Another example: A user is interfacing with the LOB data from within Excel. Roles can be used to restrict what they have access to and if the data is read-only vs read-write.

## Identity Management is Not Employee Management

I've seen quite a few Identity Management systems use the same information as the Employee/Payroll systems. While this makes it easy to shut down a user's account when an employee is terminated, it doesn't give you the flexibility you need to add users who aren't employees.

You auditor needs access to reports, but his check comes from Accounts Payable. Your vendor may need to access their sales and inventory data but that doesn't mean he's eligible for your health insurance.

It also doesn't provide you the means to provide security and access management to specific devices. There are many different interfaces these days, both public and private, for roaming and physical locations. These types of controls are often overlooked.

We all have customer web portals to access your LOB data. There are many times when you need to provide employees access to customer web portal information. Using a central identity management system, which is not HR-based, you can provide access to customer information for the employee, as well as provide customer information to the employee.

## Identity Management vs Interface/Device Management

Identity Management is more than "user credentials." Using an API-centric approach to development, you begin to see the bigger picture.

In the early days – and even now, in many shops – developers will look at the user as the primary qualifier of what data and programs are accessible. The user becomes the definer of how things are to be done. This is great for internal business applications, with the user in the office, at a desktop, with all of your intranet security in place. But what happens when the user is using the LOB software outside of this relatively uniform and protected bubble?

Do you want to provide "save as PDF" or "export data to Excel" features for accounting information when the user is at home? Do you want a user to access that data from outside at all?

Another place where this becomes important is with specific types of devices. This falls under "business discipline and procedures" more than under "security and access." Do you need to change the interface, presentation, and workflow for users accessing your LOB when in the warehouse, but give them full access when they are in the main office? What about mobile access?

What about the IoT (Internet of Things)? If you have a sensor transmitting data to your LOB, you probably don't want it interfacing with your web, GUI, or green screen. Linking the user identity as a device type of "sensor" will provide additional security because it won't have a "human" interface. Instead it might have a RESTful service or some other streamlined API.

Expanding user/password or user/password/device allows you to deny ac-

## 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**

1411 Broadway, Suite 1220

New York, NY 10018

**(212) 967-0505**

(800) 423-1964 x 302

Fax: (212)947-8593

Email: mh@eslly.com

Consultants Needed Nationwide

cess to an authentication request that doesn't match all three.

### Built-In Identity Management vs Application Identity Management

As I said at the top of the article, there is a natural tendency to use the built-in identity management of the database or O/S. This may not be the best solution.

The built-in systems are mainly used to address the security aspect of their one part of your infrastructure. They weren't designed to handle application access and security management. Nowadays, that simple login concept is too simple for restricting applications usage. It doesn't provide the flexibility you need when interfacing with your LOB application in a stateless interface, or when working with web applications, or anything non-persistent.

While I would recommend using the built-in system to interface security as a part of your process, I wouldn't recommend using it for user security. For example, if you are using an API from Excel, you can use the security to provide initial authentication and then do a secondary check using application identity management based on the user logged into Windows. If they

both pass, then you can then request the user's identity.

This allows you to implement core security features based on the interface for features like ODBC or REST services that are built-in, and then layer additional security for the user's procedures. You get the flexibility to define user accesses for web interfaces, or other stateless interfaces like EDI and web callouts while still leveraging existing functionality. This style of identity management also allows a developer to hook into other centralized authentication systems, like ActiveDirectory or OpenLDAP for role management.

In addition to external role management, it also allows the developer to interface with external login methods like biometrics (fingerprint readers and retina scanners) or single sign-on interfaces. When your company acquires another company, this sort of flexible thinking will make it easier to integrate their scheme with yours.

### Consideration for Roles Management

It is really easy to build a system that can become very complex to manage. As your LOB application becomes more customized, developers may establish role after role arbitrarily. Then, in the future, managing or remembering

what each of the roles do, and why, becomes troublesome. Discipline is key.

When a role is created, a developer should provide an explanation of what the role is used for, where it was used, and why it was created in both the role record, and within the application code. There are many times when a role will trigger some kind of override or additional access logic, and the user will need to see why they do not have access to do something. They will also need to know who has the access they need.

Another often overlooked feature of role management is remote authentication. There are many times when a supervisor needs to give access to an employee, but they are not physically present, or cannot be physically present for one reason or another.

A simple phone call to the supervisor should allow them to provide the proper authentication without giving away their password. A good way to do this is a simple mobile app that a supervisor can access to provide this authentication.

### Prototype

A simple prototype of an Application Identity Management system would be:

### USER.SECURITY

This would keep track of the user security information. Its only purpose is to authenticate the user with the LOB application. This would be a user specific settings needed for login.

Key: UserId  
<1> = Encrypted Password  
<2> = Date Create  
<3> = Time Created  
<4> = Password Expire Minutes  
<5> = Idle Expire Minutes  
<6> = Password Set Date

## IT audits have you jumping through hoops?



**PRC** can help you meet your compliance requirements and make IT more agile and productive. No extra work, nothing to remember, nothing to fall through the cracks. Our software development lifecycle tool automatically prevents or detects change according to your criteria. You can deploy, rollback, test and report quickly, automatically and with confidence. Let PRC protect your company's valuable U2 data and software assets.



SJ+ Systems Associates • [info@sjplus.com](mailto:info@sjplus.com) • <http://sjplus.com>



<7> = Password Set Time  
<8> = Last Login Date  
<9> = Last Login Time

## INTERFACE.SECURITY

Interface security is used to provide access to specific interfaces, devices, and processes. It is not specific to a user, but could be. It is designed to provide independent security options based on type of interface instead of type of user.

Key: InterfaceId (ANY, TELNET, API, EXCEL, WEB)  
<1> = Date Created  
<2> = Description of what the interface is for  
<3,n> = Inherited InterfaceId - Interface Id to use as the base settings, that this interface will be modifying.

## SECURITY.ROLE

This role security is used to define what users and interfaces have access to. Each business process, screen, or API interface may have a RoleId here. This would define if the user has access to this role, or if the interface has access to this role.

Security Roles can inherit the settings from other roles. This provides a lot of flexibility, but also provides additional level processing.

Key: RoleId  
<1> = Date Created  
<2> = Description of Role  
<3,u> = UserId  
<4,u> = UserId Date Added  
<5,u> = UserId Grant/Deny  
<6,i> = InterfaceId  
<7,i> = InterfaceId Date Added  
<8,i> = InterfaceId Grant/Deny  
<9,r> = RoleId to Inherit  
<10,r> = RoleId Grant/Deny

## API

AUTHENTICATE.USER.ID – This API would be used to authenticate a USER.ID.

Request:

- USER.ID – The USER.ID to authenticate.
- USER.PASSWORD – The USER.PASSWORD to use to authenticate.
- INTERFACE.ID – The INTERFACE.ID that is being used to authenticate.
- ROLE.ID – A ROLE.ID to check the USER.ID against.

Response:

- STATUS – Evaluates to success, failure, or expired.
- DESCRIPTION – A human readable version of the response.
- TOKEN – An authentication token that can be used for stateless and non-persistent connections.

AUTHENTICATE.USER.TOKEN – This API would be used to authenticate a token either provided by AUTHENTICATE.USER.ID or another token system.

Request:

TOKEN – The USER.ID to authenticate.

INTERFACE.ID – The INTERFACE.ID that is being used to authenticate

ROLE.ID – A ROLE.ID to check the token against.

Response:

STATUS – Evaluates to success, failure, or expired.

DESCRIPTION – A human readable version of the response.

GET.USER.ROLES – This API would be used to get a list of all the roles the user is part of.

Request:

USER.ID – The USER.ID to use in retrieving roles.

TOKEN – If USER.ID is not provided, then TOKEN should be so the application knows who made the request.

Response:

ROLES – List of all the roles (included inherited roles) for this user.

AUTHENTICATE.USER.ROLE – This API would be used to check role access for a user. This would be done after the user has logged in.

Request:

USER.ID – The User Id to get roles for

TOKEN – If userID is not provide, then TOKEN should be so the application knows who the request is for


Response:

STATUS – Evaluates to valid, invalid, not authenticated, or override request.

DESCRIPTION – A human readable version of the response. On override request, display the reason for override and who can override.

## Conclusions

Modern LOB applications aren't just about the data anymore. They are also about the APIs that the LOB provides. A developer needs to have an identity management system that provide the flexibility to handle not just data security but interface and procedures management. **IS**



**NATHAN RECTOR**  
President

International Spectrum  
nathan@intl-spectrum.com

# Flawgorithms

BY CHARLES BAROUCH

**J**ohnny is a boy. Johnny has freckles. Therefore, all boys have freckles. This is a syllogism: two true statements which are followed by a false conclusion which seems to be supported by the facts. I've met a lot of business professionals who make decisions based on logic which is even less solid. Are you running a business — or supporting one — based on flawgorithms?

## Product vs. Project

One of my pet peeves is when I buy a product and find out that I've received a project. Recently, I worked on a Drupal site. Granted, I didn't buy Drupal (in that I was not charged for the software) but I bought it in the sense of my spending time, trouble, and effort on it. Annette, my new friend, picks a template. She and I, both technical, dig in, looking for how to change the pictures on the slider. Guess what? We find instructions on how to *edit the HTML* of the module. There is no interface for selecting the pictures. Everything is hard coded.

This theme is a *project*. Software that requires you to poke at the source to make it work isn't complete. A *product* should have interfaces for the obvious needs. I accept that something esoteric — I need to handle currency fluctua-

---

*Software that requires you to poke at the source to make it work is a project. A product has interfaces for the obvious needs.*

---

tions by adding a variable conversion charge, for example — might require custom programming, but that should be the exception. Not knowing the difference between a finished product and a work-in-progress is a logical failure; a flawgorithm.

I'm sure all of us have seen this in far more critical settings. Annette simply picked another theme and we were back to work. When the project is your primary business software, that's another thing entirely. I watched a client move off of MultiValue and onto a project that was perhaps fifteen percent of what they needed. This was over a decade ago and I still recall walking into their office to the daily — sometimes more than daily — complaint: "The developer is pushing an update. I can't do any work until it finishes." Every change required a software change. Nothing was parametric.

## Go With What You Know

If you know me, you know that I never like to find myself with only one tool for a given job. I use pretty much every flavor of MultiValue. I use several brands of SQL. I'm a big proponent of using no database on certain classes of projects. Despite that, I built a product with absolutely the wrong technology recently. My flawgorithm? Go with what you know.

I love Delphi. I've built web tools with it, GUI apps, and even some basic Android stuff. When Dan Schmitt and I developed BeBackBy, I reached for Delphi. On the Android side, it works well. For the iOS side? Not so much. Long story short: I'm learning Cordova.

## Business as Usual

I recently attended a talk by David Jordan. If you don't know him, look him up. Very smart guy. I know him from our time together on the U2UG board. He tells the story of the tea kettle manufacturers who, using big data, see that they sell more black kettles than white. The flawgorithm? We need to cost-reduce white tea kettles to increase sales.

Why is that a flawgorithm? Because they are only analyzing their own data. Big data is frequently about looking at

## Is Your Membership Profile Up to Date?

International Spectrum is currently working on updating membership information for all our readers. We always like to have the most current information so that we can keep you informed of current news in the MultiValue Market.

Please go online and update your membership profile today!

[intl-spectrum.com/membership](http://intl-spectrum.com/membership)

your own. That means that any color of kettle not already in production was omitted as a sales option. It means that only their own pricing history was factored into the analysis. It leaves style and other factors out of the equation.


What was David's recommendation to these theoretic kettle makers? Step back. A tea kettle is a means to an end: They make hot water. The hot water business includes K-Cup machines, which cost more than tea kettles and have better margins. The takeaway? data can be accurate and still lead to inaccurate conclusions. Looking wider for the data – looking at your business from a wider perspective – offers more solid logic.

### Baby Carriages

In Disney's *Sleeping Beauty*, Maleficent sends her goblins to look for Aurora. Sixteen years later, one of her goblins announces that they've looked in

every cradle. That's right, they'd been using the original search parameters — she's a baby — for sixteen years.

Before we snicker at that, ask yourself how much of your business software has been unaltered for a decade or more? Are you running a business with embedded logic that hasn't been reviewed in a very long time? How many of us are running reports that tip over baby carriages? We all need to examine our systems for flawgorithms. **IS**



**CHARLES BAROUCH** is the CTO of HDWP, Inc. and the Publisher at HDWPbooks. You can read his writing in

*International Spectrum*, *Theme-Thology*, *Novo Pulp*, *PerhelionSF*, and the *Interrogative* series, which begins with *Tiago and the Masterless*.

# Open QM

taking multivalue ...  
where it has never been before

### Ladybridge Systems Ltd

17b Coldstream Lane, Hardingstone, Northampton, NN4 6DB, England

Worldwide distributor: Zumasys, 9245 Reasearch Drive, Irvine CA 92618, USA  
[www.zumasys.com](http://www.zumasys.com)

- Close compatibility with most other multivalue environments
- Easy migration process
- Maintenance-free file system for ease of use
- High quality documentation
- QMClient API for development of GUI and web applications
- Low licensing cost
- AccuTerm bundled at no additional cost
- Many unique features

[www.openqm.com](http://www.openqm.com)

# FROM THE PRESS ROOM



## Zumasys Makes the 2016 CRN Solution Provider 500 List for the 4th Time

Zumasys announced today that CRN®, a brand of The Channel Company, has named Zumasys to its 2016 Solution Provider 500 list. The SP500 list is CRN's annual ranking of the largest technology integrators, solution providers and IT consultants in North America by revenue.

"The 2016 Solution Provider 500 represent a total, combined revenue of over \$334 billion—a testament to their success in keeping pace with the rapidly changing demands of today's IT market," said Robert Faletra, CEO, The Channel Company. "This prestigious list recognizes those companies with the highest revenue and serves as a valuable industry resource for vendors seeking out top solution providers to partner with. We congratulate each of the Solution Provider 500 companies and look forward to their continued success."

### About Zumasys, Inc

Zumasys helps companies of every size transition their infrastructure and applications to the cloud. With Zumasys

cloud services, customers can easily access the latest software and hardware technologies over the Web, allowing them to focus on growing their core business instead of managing their IT infrastructures. Zumasys delivers personalized service, integrated disaster recovery and the confidence companies need to outsource the hosting of all their databases including SQL, Oracle and Pick Multi-Value systems.

In 2014, Zumasys acquired a string of MultiValue software technologies including:

AccuTerm - The leading Windows/Mobile connectivity solution for the Pick Market with more than 16,000 customers worldwide. Pete Schellenbach is now Zumasys' Director of Product Development.

MultiValue Dashboard - Originally developed by The Nerdery and re-launched in 2014, MultiValue Dashboard allows you to quickly and easily build web-based dashboards using traditional Pick programming methods.

OpenQM - Zumasys announced an Exclusive Worldwide Distribution and domestic support/maintenance arrangement for OpenQM, a cost-effective, efficient, high-performance database to run Pick applications.

jBASE - Zumasys acquired jBASE database from Temenos, based in Geneva Switzerland. Its contemporary architecture allows Pick-based applications to natively inter-

act with the underlying Windows or Unix operating system, and store data in SQL Server, Oracle and the cloud.

For more information, visit [www.zumasys.com](http://www.zumasys.com) ■



## Linear Hash Reporting and Resizing Tool for OI v9.4 released

In an effort to help assist OpenInsight system administrators manage Linear Hash tables Revelation Software is publishing a reporting and resizing tool called RTI\_LH\_STATISTICS. This tool will enable OpenInsight system administrators to view linear hash statistics within an application and identify and resize poorly hashed tables.

Poorly hashed/sized tables are one of the most significant reasons for degradation in system performance.

The Remake table option of RTI\_LH\_STATISTICS utilizes a program called RTI\_REMAKETABLE. RTI\_REMAKETABLE works by sampling 10% of the rows in a table, calculating the average record size, then making a new pre-sized table with the frame size larger than the average record size and a threshold of 50% to force the table to resize more frequently. RTI\_

REMAKETABLE copies the rows from the existing table to the new table, then edits the REVMEDIA record to point to the new table rather than the old and attaches the table again.

For poorly hashed tables that have been resized the .LK will be larger in size and the .OV smaller in size. The total size of .LK + .OV will most likely increase. By making a larger .LK and a frame size larger than the average record, the resize program will force most records into the primary (LK) frames. The result will be that the performance of reads and writes will increase significantly.

You can use this tool on any table, including ! and DICT files, as well as system files such as SYSLISTS and SYS-OBJ. Please be aware that the underlying REV number will change, for example as table whose DOS name was REV12345.LK may become REV12890.lk. You can run the window while the system is busy, but when resizing tables all users and background processes should be stopped. The process will update the DBT to reflect the new DOS filenames.

RTI\_LH\_Statistics is available at no cost to all current Works members and can be downloaded from the Works Download section.

### About Revelation Software

Founded in 1982, Revelation Software delivers a suite

# FROM THE PRESS ROOM

of application development tools and companion services that take full advantage of leading network computing architectures, messaging, groupware, and client server platforms. Today, the company's flagship product OpenInsight is the only database development and application environment that provides both Windows and Java-based GUI s tools to develop and deploy web-based and client server applications that support native and relational XML, SQL, Lotus Notes and the leading legacy Multi-Value data sources such as Arev, Pick and IBM Universe. There are more than 1.5 million licensed users of Revelation products across 60,000 deployed sites worldwide. The company has offices in Westwood, New Jersey, as well as a European distributor in the United Kingdom, and an Asia Pacific subsidiary in Australia.

For more information, visit [www.revelation.com](http://www.revelation.com) ■



## Rocket MVS Toolkit 2.2 Now Available

MVS Toolkit 2.2 includes numerous enhancements and fixes that will help developers more easily work with the Toolkit. Examples of these enhancements are as follows:

Support for export and import of a web service operation within a web service , which simplifies deployment of updates by eliminating the need to export and import an entire web service when the web service operation changes.

Support for cross origin resource sharing . We have added a cross origin resource sharing protocol that enables a web server to serve resources to web based clients that access resources from multiple domains. Some web APIs enforce a "same domain" policy, where the client can only access resources from the same domain. This protocol provides a means of allowing access to resources on other domains.

Improved performance when a web service returns a large result set.

Support for nested web service URLs.

Find the Readme file for MVS Toolkit 2.2, including upgrade requirements, licensing, other considerations, known issues, and enhancements / resolutions at:

[www3.rocketsoftware.com/rocketd3/support/documentation/mvs/Readme.txt](http://www3.rocketsoftware.com/rocketd3/support/documentation/mvs/Readme.txt)

### About Rocket Software

Rocket Software is a global software development firm that builds and services Enterprise Infrastructure products for the world's leading

OEMs, networks and software companies and enterprises. The company's current lines of business complement and extend strategic OEM offerings in the areas of database, business intelligence, storage, networks and telecom, terminal emulation and FTP, integration, modernization and SOA, and security. Rocket is engaged in business and technology partnerships with IBM, EMC/RSA, HP/EDS, Nortel, Motorola, and many others. Rocket Software is based in Newton, Massachusetts.

For more information, visit [www.rocketsoftware.com](http://www.rocketsoftware.com) ■



## Kourier Integrator Release 4.3 now Available.

Kore Technologies announces the latest version of its award-winning enterprise integration and data management suite. Having successfully completed beta-testing over the last several months, Kourier Integrator Release 4.3 is now generally available.

The new release of Kourier Integrator simplifies the development and usability of RESTful Web Services and introduces the Kourier REST Gateway to provide secure data access to MultiValue ap-

plications from outside the corporate firewall.

The Kourier REST Gateway is a critical component of the real-time integration architecture. It's responsible for providing secure access to MultiValue applications from outside the firewall while it monitors, manages and measures REST Resource usage. Organizations deploying the REST Gateway can feel confident about offering access to their system from the outside world without worrying about a user consuming all of their system resources using Policies to limit (rate) the maximum number of requests per minute/hour/day for each Gateway user. An interactive dashboard offers visualization of REST API request trends and performance, giving administrators actionable data for making system adjustments.

Kourier's REST development framework has been enhanced to streamline and simplify how developers build real-time REST APIs. Developers will be more productive because they can focus on the application interface instead of low-level protocol details such as data validation, resource security and transaction logging. Creating REST Resources is done predominantly via configuration pages instead of programming. For more advanced requirements, developers can use Event Handlers to leverage existing application business logic or add special instructions at specific tim-

# FROM THE PRESS ROOM

ing intervals. API consumers also benefit from using Kourier REST APIs because they are more developer-friendly. REST features such as pagination, field limiting, dynamic sorting and query wildcards can be utilized with no additional programming effort, reducing the time and cost required for application integration

Release 4.3 also includes significant improvements to Kourier's support for additional data sources (e.g., Oracle, Microsoft SQL Server, PostgreSQL) with the new Quick Start Workbench and advanced data profiling system. Users can quickly create a near real-time Microsoft SQL Server data warehouse using the profiled data to generate export definitions and maintenance-free Microsoft SSIS packages.

## About Kore Technologies

Kore Technologies is a leading provider of enterprise integration, business intelligence and eCommerce Web solutions for MultiValue and Microsoft SQL Server databases. Founded in 1999 as a professional services and software development company, Kore specializes in providing business-to-business connectivity and best-in-class product integration with its flagship product—Kourier Integrator. Leveraging this technology, and Kore's extensive experience with enterprise applications and database systems, Kore delivers solutions for integrating disparate

applications, building data warehouses, and developing business intelligence solutions.

For more information, visit [www.koretech.com](http://www.koretech.com) ■



## Ladybridge Releases 3.4-6 of QM

This release is mainly to correct an issue described in section 5 but also includes the following enhancements.

The DECRYPT.FILE command can be used to cancel encryption for a file or specific fields.

The OSCOPY command and corresponding !OSCOPY() subroutine have been extended to allow appending to an existing file.

The QMBasic ELEMENT.EXISTS() function can be used to determine whether a named element is present in a data collection.

The compiler no longer reports an error if a subroutine argument passed as a matrix is not dimensioned.

The QMSAVE tool (see <http://www.openqm.com/cgi/lbscgi.exe?T0=h&t1=kb.00011>) has been modified to add verification of the file and record counts.

## About Ladybridge Systems Ltd

Ladybridge Systems was founded in 1992 by former Prime Computer personnel and is headquartered in Hardingstone, Northampton, England. The primary activity of the company is the OpenQM multivalued database for which Zumasys were appointed as worldwide sales and marketing partner in January 2015. Ladybridge Systems also provide training, consultancy, development, and maintenance services to MultiValue databases in the UK and worldwide. Ladybridge Systems remains a small privately owned company, offering a personal level of service to clients around the world.

For more information, visit [www.ladybridge.com](http://www.ladybridge.com) ■



## Entrinsik Opens European Office and Strengthens Global Presence in the Business Intelligence Market

Entrinsik Inc. has expanded its international presence with the opening of its European office.

Entrinsik, a global organization headquartered in North Carolina USA, has been de-

veloping software for over 30 years providing high quality, innovative data management solutions to businesses and organizations across various industries.

Entrinsik Informer business intelligence software enables organizations to perform ad-hoc reporting and analysis, blending data from multiple sources to create interactive reports and visualizations. Regardless of technical expertise, users can easily customize views, drill down into underlying data, dynamically group & sort, export & embed, all in real time, all on one screen. Informer is used by over 1500 organizations around the world in various industries, and is top-ranked by various international research organizations including BARC.

Entrinsik opened an office in Europe with the help of Eurodev ( [www.eurodev.com](http://www.eurodev.com) ) a Netherlands-based business development company. The goal is to deliver Entrinsik products to the European market to meet the growing need for innovative reporting and data analysis solutions. With Informer, organizations across Europe can access key information, identify trends, uncover opportunities, and make data-driven decisions with confidence.

This global expansion will allow Entrinsik to better support its existing client base in Europe, while continuing to aggressively grow the business internationally. The company

# FROM THE PRESS ROOM

currently works with several distributors of its Informer software in Europe, Australia, and South America, and with the opening of Entrinsik Europe the company is building a team to consolidate all European sales and marketing efforts locally.

"With strong credentials across multiple industries and an international client base, Entrinsik can now better provide European organizations with the tools necessary to help them succeed and stay competitive in a highly dynamic business environment," says Doug Leupen, President/CEO at Entrinsik. "Eurodev's business development expertise has proven invaluable."

Entrinsik's European office can be reached by calling +31 548 659 026 or by email at [europa@entrinsik.eu](mailto:europa@entrinsik.eu).

## About Entrinsik

Entrinsik, Inc. is the market leader in providing web-based operational and analytical reporting tools to the MultiValue marketplace. Today, with Entrinsik Informer, thousands of end-users using MultiValue databases enjoy the ability to leverage the Web to share information assets with virtually unlimited number of people anywhere at any time. With its intuitive, easy-to-use interface, Informer provides end-users, IT staff, and executives a business intelligence environment and consistent, real-time information impacting key business

processes. Spectrum Magazine has described Informer as a "supercharged reporting tool...designed for decision-makers."

Founded in 1984 and in the MultiValue market since, Entrinsik aims to provide leading edge technology and unsurpassed service to their customers.

For more information, visit [www.entrinsik.com](http://www.entrinsik.com) ■



## Drexel Management adds EMV support for Ingenico iSC250

Drexel Management adds support to CCListenDSS for the Ingenico iSC250. This expands on the existing support for the Verifone VX805 hardware.

Most financial institutions and card issuers are sending their customers chip cards that require signatures for verification, rather than PINs. While both types of cards improve security, PINs protect against forged signatures and are more secure. Most major retailers will be equipped to accept PINs or signatures and are urging card issuers to offer customers chip-and-PIN technology.

In addition to the EMV Card support, NFC (Near Field Communications) technology has been added to support ApplyPay transactions. Drexel Management's CCListen is a database agnostic middleware that support seamless integration between any MultiValue and Non-MultiValue application. CCListen includes support for credit card and debit card processing, card vaulting, Level 1 and Level 2 card processing, as well as tokenization to fully conform to PCI requirements, in addition to US Pin Debit and store gift card programs. Partner programs are available as well.

## About Drexel Management Service

Drexel Management Services specialize in Windows based products, enhancements to existing MultiValue based applications, and customized solution based programming:

- Raining Data D3 Pick database conversion and/or upgrade
- JBase installation, conversion and support
- Custom .Net Applications
- Networking
- VB Custom Applications
- VPN Internet Connections
- Web Applications

There are many other types of core services and products that we, Drexel Management Services, specialize in.

For more information, visit [www.drexelmgmt.com](http://www.drexelmgmt.com) ■



## Carnation Software releases MacWise Version 15.2.7

Carnation Software releases MacWise Version 15.2.7 that fixes intermittent hangs when switching between windows. It would only affect you if you have more than one MacWise window open for multiple sessions.

Other recent updates added support for special number characters typed on the keyboard that use the Option Key. Also, support for System Builder commands to launch web pages or files.

Upgrades are free for users who have an earlier version 15. Exceptions apply for site licenses.

NOTE: Version 15.2.7 has been tested with the beta version of the next Mac OS (Sierra) and no problems have been found. Will be testing each beta version of the OS as it is released but it looks good so far.

You can see all of the version history here...

[www.macwise.com/MacWiseXVersions.html](http://www.macwise.com/MacWiseXVersions.html) ■

# QM and VFS

BY BRIAN LEACH

I had a chance last month to look at an aspect of the QM platform that I had not used before: The VFS or Virtual File System.

The background to this was an interesting request from a client to look into various options for integrating their application with SQL Server. Not just as a reporting base or data warehouse, which is bread and butter stuff, but to host their live data in SQL Server and have their application make use of it for transactional purposes. Having pointed out the obvious risks and performance issues to be expected by going that route — given that they have a very large, complex, mature and hard working UniVerse application managing huge numbers of complicated transactions and talking to numerous upstream and downstream systems — and the fact that a MultiValue application is designed in a fundamentally different way to an SQL one — I duly set out to assess the current state of various technologies that might allow this.

The actual findings are commercially sensitive. I'm not going to reproduce those here. Nor am I going to list out all the various options that I put before them, or recommend one option over another. The goal was to lay out alternatives so the client could decide

---

*The QM VFS is a means of redirecting file access. It works on a file by file basis, so that you can mix and match VFS files with regular files...*

---

which ones would merit further consideration.

I knew that QM had a virtual file system, but I had never had a reason to play with it. It's also a fairly undocumented experience, with only a couple of sample drivers to look at, neither of which was a good model for what I was attempting. I needed to develop a proof of concept driver that would talk — at reasonable speed — to SQL Server.

Since this was under the auspices of producing a report, development time and budget would be extremely limited. But I wanted to see whether it was technically feasible to create something that might perform similarly to, for example, MVON. I knew the experience might not be as smooth and streamlined as the On Group's approach to running MultiValue on SQL. The end result was likely to be more of a hybrid approach than MVON's seamless model.

As you can tell, I had already been highly impressed with the way that MVON 'just worked' straight out of the box and how quickly I had been able to get a demo system running on MVON. So, I knew that was a high benchmark to aim at.

## Using VFS

The QM VFS is a means of redirecting file access. It works on a file by file basis, allowing you to mix and match VFS files with regular files. You can, for example, host a data file in VFS and the associated dictionary as a native QM file.

To access a file using VFS you need three things:

- A VFS driver.
- A VFS server entry.
- A modified file pointer.

The VFS driver is the fun part but I'll speak to that later in this article. For now it's just worth saying that there are two types: (A) Internal drivers written as QM Basic classes and (B) external drivers written in native code. For various reasons my driver is external. I did also model an internal driver but found it too limiting. In other situations the mileage may vary.

A VFS server is a logical collection of information that defines a particular data source. It combines the name of



## 10 Years Strong & Growing

Integrity, Efficiency, Service,  
Guarantee, Global Productivity,  
Midwest Rates, Universe,  
UniData, D3, Barcode Services,  
Nationwide Support, Web Services,  
Custom Barcode Programming,  
.Net Programming, Mentoring,  
Project Outsourcing



[www.pickprogram.com](http://www.pickprogram.com)  
[contact@pickprogram.com](mailto:contact@pickprogram.com)  
(614) 921-9840

**Do you... have projects sitting on "the list" and not being completed?**

**Do you... just need some assistance from time to time?**

**Do you... want up-front, guaranteed estimates?**

**Senior Consultants are Currently Available.**

*"PICK Programmer's Shop has been outstanding to work with. They work with the highest integrity, are very responsive and quite knowledgeable. I can't imagine operating without them."*

*Patti Rowlette, Rowlette Executive Search*

***It's not just an anniversary. It's a guarantee.***

the VFS driver with connection and host details, so you could have a generic driver called, say, demovfs (for the want of anything original), and it could talk to a number of different databases or external systems for pulling

and pushing data. It is similar in that regard to an ODBC data source definition. The VFS server definition is created using the SET.VFS.SERVER command, which is necessary as it encrypts the password for storage.

To add a VFS server called 'DEMO' with an external VFS driver of 'demovfs' accessing a remote data source named 'MY\_DBSERVER' you would need the SET.VFS.SERVER command as seen in [Figure 1].

The password is stored, encrypted, and passed to the driver since you don't want it to prompt you for authentication information when it is actually running. It could be a dummy, as the driver can decide exactly what it wants to do with the parameters it is given.

The LIST.SERVERS command lists your available VFS servers [Figure 2].

Once you have defined your server, actually making use of it is satisfyingly simple. Create a file pointer with a path name composed of the following:

```
VFS:server_name:target_name
```

The target\_name is just a string that just gets passed through verbatim to the driver to identify the final data source. In my case I'm using the name of the SQL table or view, but it could be absolutely anything just so long as

it is meaningful to the driver. You might, for example, create an IMAP driver that could interrogate a mail box passing the name of a folder or tag.

So to create a VFS file called MY\_FILE that will instruct the DEMO server

to access a thing called MY\_TABLE (you get the naming convention), you would create a QM file as follows. Note that I'm still defining a local dictionary for the file [Figure 3].

You can now access MY\_FILE just like any other QM file: listing it, reading and writing records and generally having fun. All of these actions just

```
SET.VFS.SERVER DEMO EXT demovfs MY_DBSERVER my_username my_password
```

Fig. 1

```
:LIST.SERVERS
VFS Servers..... IP address..... Port... Sec Remote
user name....
DEMO                Handler: EXT demovfs
                    MY_DBSERVER                0           No      demo
```

Fig. 2

```
:CREATE.FILE DICT MY_FILE
Created DICT part as MY_FILE.DIC
Added default '@ID' record to dictionary
:ED VOC MY_FILE
VOC MY_FILE
3 line(s)
----:
0001: F
----:
0002:
----: r VFS:DEMO:MY_TABLE
0002: VFS:DEMO:MY_TABLE
----: fi
'MY_FILE' filed in VOC
```

Fig. 3

get passed to the driver via the VFS API and it can decide what to do with them.

### The VFS API

The nice thing about VFS is the API itself. The nice thing about the API is what it leaves out.

The VFS API is really, really simple, and makes absolutely no assumptions about what you want to do with your virtual files under the covers. Where other approaches might, for example, use a mapping schema that resolves into generic SQL statements being passed to an ODBC-like driver, VFS works at the logical MultiValue level and lets you get on with the job of doing your own mapping underneath. It knows that you want to open a file, read, write and delete some records, possibly clear the file from time to time and run selections, but how you do that is entirely up to you.

Now there is no doubt that if you want to access an SQL database and haven't got the knowledge (or time, or interest) to set up all the mapping and connection details yourself, having a mapping schema is a great benefit. But it also restricts you to whatever that schema will allow. If the schema is determined to turn your writes into generic UPDATE, INSERT and DELETE commands you don't have the choice to use alternatives such as SQL Server specific MERGE or TRUNCATE statements.

If you have a tame TSQL guru who is willing to churn out highly optimized stored procedures, or even to make use of the CLR features inside SQL Server, then again you need a tool that will allow you to make full use of those.

### Mapping the Data

In my case this was only a proof of concept, so I created a simple driver in C# that expected each target to be either a table or a view. The SQL Server .NET client exposes a schema API through which you can easily discover the layout of any table or view. I wanted to test two layout options: a simplified storage consisting of just a primary key and a varchar(MAX) field to hold a record body just as it would appear in a MultiValue file; and a normalized layout consisting of columns of different data types that would correspond one to one with fields in the record by their ordinal position. The only requirement was that both layouts would require a primary key, again discoverable through the schema API. The driver would be responsible for translating between the MultiValue representation and the columns in the table or view, and would handle some internal niceties as turning PICK format dates into SQL Server equivalents [Figure 4].

### Merging Data

Writing to an SQL database is a more painful operation than writing to a

MultiValue database, not the least because of the differentiation between an insert and an update. Often the combination is termed an *upsert* operation, which is normally handled by either (A) testing explicitly for the presence of a row with the given primary key *before* deciding whether to insert or update; or (B) by attempting an update and following this with an insert if it fails. Both of these are tried and tested methods and work across different database platforms.

The MERGE statement in SQL Server combines these into one and can be quicker in some circumstances, though with a verbose and ugly syntax. It is normally used for merging the content of two tables, but your statement can just select the input parameters as the merge source.

Also, like all SQL Server statements, it needs to be planned and optimized and so only comes into its own through reuse. SQL Server will cache and reuse the latest query plan where possible, but to ensure a level of performance you almost always need a stored procedure: so the driver discovers or creates such a procedure on the first write request it receives to a given VFS file. If the file is only ever read, no procedure is generated. Further, the driver caches the parameterized command objects so they are there ready for use and it doesn't have to waste time checking the parameters again. These are all small things, but the cumulative performance benefits can stack up.

So having created (and unit tested!) the generic SQL Server actions, it is time to turn to the VFS part.

### The VFS Driver

As I wrote above, a VFS driver can be implemented in two ways: as an exter-




Fig. 4

nal driver written in a native language such as C or C#, or as an internal QM Basic class. Both internal and external drivers offer the same API but with one important difference. For readers more familiar with other MultiValue databases, QM Basic has object oriented features similar to VB or VB.NET that sadly the other manufacturers have never picked up - the local functions in U2 are a poor alternative. QM classes are written in QM Basic and surface methods as local subroutines and functions.

The internal driver, the one written in QM Basic, is much simpler to develop. There is a template class that you can copy and from there you can simply fill in the methods provided: they all have reasonable comments and the parameters are meaningful. However, since QM Basic doesn't talk directly to SQL Server in the way that I wanted, I could only work the internal driver via an intermediate service application over a socket connection. Also, a new instance of the class is created for each separate VFS file being accessed: and the driver persists for as long as the VFS file is open, which means that

unless the file is opened to a named common the driver is constantly being created and destroyed. For a staged architecture, this made it very inefficient as it had to continually reinitialize to reestablish its connection with the service.

The external driver, on the other hand, could be written as .NET executable (or any other language that can communicate through the API) and so did not need an intervening service to talk to SQL Server. More importantly, it is launched as a single instance for each logged on user accessing files through a VFS server definition. Instead of associating each VFS file with a separate instance of the driver, and the internal driver does, a file identifier is simply passed to or from the API on each call and the same driver manages a whole group of files. This meant that the driver could itself maintain a single connection to SQL Server for that user for lifespan of their login session.

The first request to open a VFS file causes QM to create the driver by starting the executable and passing to it the identifier of a named pipe. This can act like a TCP network socket, but

will also route via shared memory for speed if both end points are on the same machine. The executable must open its end of the pipe and can then receive commands and send responses through the pipe to the QM session. One complication is that under .NET there is support for named pipes, but only in a specific configuration that does not match the way that QM uses them, and so a driver written using .NET needs to invoke the underlying Windows API to establish the connection [Figure 5].

A minimal driver needs to handle the following actions, which are identified by an action code passed at the start of each request and a set of parameters that are passed up and down the pipe in a simple encoded fashion [Figure 6].

Additional functions allow you to chunk large selections so that you do not pass huge lists of keys back through the pipe in a single hit; to work directly with indices, and to lock and unlock the file.

### Try it for Yourself

I can't share the code as it's creation was partly funded by the client for this report and so it is not mine to give. But hopefully this article provides sufficient clues to allow you to write your own: it could be an interesting side project.

There are a few niceties in regard to formatting the data returned that I will leave up to the reader to discover - I wouldn't want to deny the reader too much of the fun! **IS**

**BRIAN LEACH** is a MultiValue developer, consultant, trainer and author, and a board member of the Rocket Software Users Group. You probably knew that. Find him at <http://www.brianleach.co.uk>.




Fig. 5

Action	Description
V_SHUTDOWN	Close the driver
V_CLOSE	Close the file
V_FILEINFO	Provide details on the file and its indices
V_READ[V]	Read a record or field
V_SELECT	Create a select list of keys


Fig. 6

# MARKETPLACE

## ACCOUNTING


**Natec Systems**

[www.natecsystems.com](http://www.natecsystems.com) | [nrector@natecsystems.com](mailto:nrector@natecsystems.com)




### QuickBooks API for the MultiValue Database

- Read/Write Directly to Quickbooks Databases
- mvQB API is Designed for the MultiValue Program to Use
- No Need to Learn the Internals of QuickBooks
- QuickBooks Pro/Premier/Enterprise



**NATEC Systems**  
Providing Solutions to your MultiValue Questions



Phone: 303.465.9616  
E-mail: [mvqb@natecsystems.com](mailto:mvqb@natecsystems.com)  
Website: [www.natecsystems.com](http://www.natecsystems.com)

## COMPLIANCE

**SJ+ Systems Associates**

[www.sjplus.com](http://www.sjplus.com) | [sjoslyn@sjplus.com](mailto:sjoslyn@sjplus.com)

## CONSULTING

**Drexel Management Service**

[www.drexelmgmt.com](http://www.drexelmgmt.com) | [dconboy@drexelmgmt.com](mailto:dconboy@drexelmgmt.com)

**Execu-Sys, LTD**

[www.eslny.com](http://www.eslny.com) | [mh@eslny.com](mailto:mh@eslny.com)

**HDWP**

[www.HDWP.com](http://www.HDWP.com) | [results@HDWP.com](mailto:results@HDWP.com)

**Modern MultiValue, LLC**

[www.ModernMultiValue.com](http://www.ModernMultiValue.com) | [info@ModernMultiValue.com](mailto:info@ModernMultiValue.com)

**PICK Programmers Shop**

[www.pickprogram.com](http://www.pickprogram.com) | [brian@pickprogram.com](mailto:brian@pickprogram.com)

**Precision Solutions**

[www.precisiononline.com](http://www.precisiononline.com) | [Kevin@PrecisOnline.com](mailto:Kevin@PrecisOnline.com)

## DATABASE

**Ladybridge Systems Ltd**

[www.ladybridge.com](http://www.ladybridge.com) | [sales@Ladybridge.com](mailto:sales@Ladybridge.com)

## REPORTING

**Brian Leach Consulting, LTD**

[www.brianleach.co.uk](http://www.brianleach.co.uk) | [brian@brianleach.co.uk](mailto:brian@brianleach.co.uk)

## TERMINAL EMULATOR

**Zumasys**

<http://www.zumasys.com/products/accuterm/>



AccuTerm® software, the leader in terminal emulation, allows you to access your MultiValue application—whether it is on-premises or in the cloud—from any Windows device.

[www.zumasys.com/accuterm](http://www.zumasys.com/accuterm)  
818-951-1891



## LETTERS TO THE EDITOR

Have an opinion on an article: Agree, disagree, or enhancement to an article from a previous issue? International Spectrum and our authors are interested in hearing from you!

E-mail: [editor@intl-spectrum.com](mailto:editor@intl-spectrum.com)

### WANT TO SEE A SPECIFIC TOPIC?

International Spectrum is looking for writers, feedback, and topic ideas. We all have specific topics and issues that we need answers to find solutions for. Send us an E-mail with topics you would like to have covered in the magazine or on the website.

E-mail: [nathan@intl-spectrum.com](mailto:nathan@intl-spectrum.com)

### WANT TO WRITE?

Expand your professional credentials, and provide us with an article.

Give us a rough and ugly outline, and we will help you refine it, proof it, and make it press ready. Or you can give us something polished, proofed, and press ready to publish.

Share your thoughts and expertise with over 10,000 fellow MultiValue developers and users.

E-mail: [editor@intl-spectrum.com](mailto:editor@intl-spectrum.com)

### NEED A MENTOR?

Mentors give developers the ability to ask industry experts for direction, code examples, and/or just ask them to see if something makes sense. Sometimes, all you need is a resource or example to start or complete a project.

Check with us to see who is available for mentoring, and how you can take advantage of it to save your business or company money.

E-mail: [nathan@intl-spectrum.com](mailto:nathan@intl-spectrum.com)

### WANT TO BE A MENTOR?

We have many retired or semi-retired professionals out there that would love to share their knowledge of MultiValue development. If you are one of them, please contact us to see what mentoring is all about.

E-mail: [nathan@intl-spectrum.com](mailto:nathan@intl-spectrum.com)

## IS.HASH.SHA1

### Generating with UniBASIC DIGEST Command

In cryptography, SHA-1 (Secure Hash Algorithm 1) is a cryptographic hash function designed by the United States National Security Agency and is a U.S. Federal Information Processing Standard published by the United States NIST. SHA1 is used in many places to generate a unique hash value representing a string or file. It is widely used in place of MD4 and MD5 hash as a more secure hash value. While not as secure as SHA256, it is still used in many places for data integrity, version control, and other

features that need an unique one-way signatures.

#### SHA1 Function

SHA-1 produces a 160-bit (20-byte) hash value known as a message digest. A SHA-1 hash value is typically rendered as a hexadecimal number, 40 digits long.

```
SHA1("The quick brown fox  
jumps over the lazy dog")
```

gives hexadecimal:

```
2fd4e1c67a2d28fced849ee1bb  
76e7391b93eb12
```

SHA1 will detect even small changes in the string and cause the returned hash value to change. You can always read more about SHA1 at:

<https://en.wikipedia.org/wiki/SHA-1>

#### UniBASIC Hashing

UniBASIC comes with hashing extensions in the form of DIGEST(). The DIGEST has SHA1 built-in. This allows you to use code like the example in Figure 1

The subroutine found with this article is based on the code above, but has been structured to be interchangeable with non-UniBASIC versions of the SHA1 subroutine also found on the International Spectrum Website at: <http://www.intl-spectrum.com/resource/756/default.aspx>.

#### Example

See Figure 2.

#### Output

See Figure 3. IS

```
SHA1.MESSAGE = "The quick brown fox jumps over the lazy dog"  
ret = DIGEST('SHA1',SHA1.MESSAGE, 1, SHA1.HASH)  
CRT SHA1.HASH : ` = 2fd4e1c67a2d28fced849ee1bb76e7391b93eb12
```

Fig. 1

```
SHA1.MESSAGE = "The quick brown fox jumps over the lazy dog"  
CALL IS.HASH.SHA1(SHA1.MESSAGE,HASH.VALUE)  
*  
TEST.VALUE = "2fd4e1c67a2d28fced849ee1bb76e7391b93eb12"  
CRT HASH.VALUE : " = " : TEST.VALUE : " - "  
IF (HASH.VALUE EQ TEST.VALUE) THEN  
  CRT "Ok"  
END ELSE  
  CRT "Failed"  
END
```

Fig. 2

```
2fd4e1c67a2d28fced849ee1bb76e7391b93eb12 = 2fd4e1c67a2d28fced849ee1bb76e7391b93eb12 -  
Ok
```

Fig. 3

# Unstructured Data

BY NATHAN RECTOR

In the world of data storage and retrieval, life used to be simple. You had elements of data. Data had various types. Data elements went in columns in a table. And those of us in MultiValue had it even easier because we weren't even bound by those rules. But now we have to deal with unstructured data.

Unstructured data is a generic term for describing any information that is not in a database. Unstructured data can be textual or non-textual. Textual unstructured data is generated in media like email, PowerPoint, Word, instant message, and social media. Non-textual unstructured data is generated in media like images, audio files, and videos.

If left unmanaged, unstructured data can pose a costly liability when the information cannot be located, or a versioning issue between the database, or just between multiple copies of the same file.

As our users become more tech literate, they expect to be able to use their skills to work with information in the tools they know. This introduces the version control issues I just mentioned and the User Expectations vs. IT Governance controls and policies.

---

*The most overlooked unstructured data we have in the enterprise are emails, documents, faxes, images, and videos. In the good old days, all we really worried about was scanned images or cold stored printouts.*

---

Let's face it, the user doesn't know and doesn't care that you have to report reliable and audited data to someone else.

These are some of the important issues of trying to get a handle on unstructured data. But deciding how to work with it, manage it, and parse the important data points from it is just as important.

## **Management of Unstructured Data**

When working with unstructured data, and the files that contain the information, the natural reaction for a system administrator is to create an index. But the index is just one part of working with unstructured data.

The real trick is figuring out what is important about the unstructured data,

and how, when, and why we have the unstructured data. This will help you decide how to index and what to do with "data fluff."

We know from our existing databases that indexes will evolve over time as users want to interact with the data differently. Traditionally, we use these indexes to speed up access to the data.

In this case, we are not talking about indexes in the sense of speeding up data retrieval, but indexing as in categorizing, exploring, and relating the unstructured data to existing structured data. Without this relationship, the unstructured data is just data fluff.

What is data fluff? Data fluff is data that is stored because the designer or developer thought it would be important, but has no real value or use to the business process. Every database has data fluff. Sometimes it is just transitory data to indicate flags needed later or during a record correction, or audit information like date and time created, user who created it, an so forth.

Many log files can be considered data fluff because they are used only for debugging and not actually used to help define, isolate, or refine important business processes.

Don't get me wrong, data fluff can be important and useful information, but if it is not being used, it is just fluff.

### Overlooked Unstructured Data

The most overlooked unstructured data we have in the enterprise are emails, documents, faxes, images, and videos. In the good old days, all we really worried about was scanned images or cold stored printouts.

Now scanned images are only part of the problem we have to address. These days, emails are the biggest unstructured data source that people are ignoring. Chat and social media come after that, but they are all basically the same types of unstructured data as email.

If someone wants to manage their own email, they build a structure of folders that fits that person's need at the time. Fast forward a few years, and the struc-

ture doesn't handle the needs of that user, and the information they have indexed is unavailable to others who may need it.

The classic example of trying to relate this unstructured data to a structured data source is a CRM. A true CRM is designed to present anyone with relevant information about a customer. The keyword here is "anyone," not just the person directly involved with that customer.

What happens is the relevant information for that customer is in a salesperson's email or voice mail, but the email and voicemail systems are disconnected from our enterprise data. If someone in production needs access to information that was emailed to a salesperson, it is not available. Hence the development of the CRM, or the continued search

for an adequate CRM that handles all aspects of the enterprise.

Another overlooked data source is our Windows files systems. Everyone is using Excel and Word these days for something; printing labels, creating warranty documents, and data mining statistical data. Some of these files are duplicates or versions of the structured database data. Some are not, but are relevant to some aspect of the company. Even if it's just to the person who created it, they created for a reason.

### What to do with Unstructured Data

We all can see the value of indexing emails and documents sent to customers and having it available to access from our CRMs. This is the indexing aspect of working with unstructured data — creating a relationship between

# THEME-THOLOGY: INVASION

Voices by Lisa A. Kramer  
I Was a Teenage Alien by LJ Cohen  
Singularity by Jeremiah Lewis  
Not Like Us by Mike Reeves-McMillan  
That Kind by Charles Barouch  
Yellow by Bill Ries-Knight  
An Invasion of Ideas by Jeremy Lichtman  
Famine, with Fries by Jefferson Smith  
The Several Monsters of Sainte-Sara-la-Noire by Michael Williams  
Going Viral by Rachel Desilets  
Dead Planet Scrolls by Timothy Hurley  
Red Vapor by Michaela Susanne  
The Worms Crawl In by Michelle Mogil  
Nano Nation by CM Stewart  
The Woods, The Cellar, and Cover art by Aaron Wood  
All other interior art by Juan Ochoa

WWW.THEME-THOLOGY.COM



the unstructured data source and the database's structured data.

Now what do you do with it? Our production systems all have enterprise alerts for various things, like low quantities, out of date prices, and upcoming order statistics.

Let's look at some alerts and data you may be able to generate using email and instant message history. You can use this information with a predictive modeling API (Google has one) that would tell you if the mood, or sentiment, of the message is good, bad, or ugly.

Wouldn't be nice to know if the communications between your people and the customer is positive or negative or indifferent? Or if the message that the user just received showed a critical issue or negative message that needs to be addressed?

Another thing you can do with unstructured data is to map trends from your web logs — which customers visit, which ones don't.

### **Unstructured MultiValue Data**

Let's back up a few steps and look at our existing database and the data we have collected.

While we generally assume our MultiValue databases are structured data stores, we also likely have tons of unstructured data as well. For example, product descriptions, notes, and text blobs associated with customers and vendors. Much of this information can provide useful valuable data, over and above the structured data it is associated with.

Also add into the design mix the fact that MultiValue databases are not inherently structured data sources.

Yes, we do create dictionaries to impose a structure, but we are not required to follow that structure when storing information. This feature gives the old referential integrity proponents fits.

Now add into the mix the correlatives, I-types, and virtual fields that allow us to create relationships to data in other locations, manipulate databases in the same location, or even parse, extract, and transform unstructured data blobs into structured contents. You now have a database with a slight identity crisis.

### **Parsing, Indexing, and Imposing Structure to Non-MultiValue Data**

The type of unstructured data you are working with will dictate the structure you need to — or want — to impose on it.

Let's say we have a web log, or analytic file, and we want to see the relationship

between sales, time of day, and how many phone calls you have received from a customer. Ok, I realize that this data is not exactly unstructured since it is a formatted log, but that is where you and the rest of the world differ. Parsing, extracting, and transforming information is the nature of our business, so this doesn't seem out of the ordinary.

To other people, this data is totally unstructured and takes time to manipulate. Or they just do it by hand.

Let's look at a totally different type of unstructured data: emails, chat logs, CRM notes, and social media posts. All this data is in textual blobs. Extracting meaning, issues, or just searching for a specific conversation is important to your business.

While many unstructured documents, like Word, Excel, or video files, look like they are unstructured, many really are not. Many of the files you want to create relationships with have metadata storage features in them. Yes, even images and video files have metadata storage.

What is metadata storage? It is the ability to enter information about the rest of the file or blob, like author, customer, dates and time, and other structured information about the document. This provides links and relationships that you can access or populate from your structured database.

### **Summary**

This article should have provided you a starting point to think about how to work with unstructured data. Keep watch on the International Spectrum website and other magazine articles that talk about the nuts and bolts of access, parsing, and indexing various unstructured data. **IS**

The logo for mvTest features the text "Quality Counts" in a large, bold, black font at the top. Below it are two stacked rectangular boxes: a green one with the word "PASS" in white and a red one with the word "FAIL" in white. At the bottom, the text "mvTest" is written in a large, bold, dark red font, with the URL "http://www.brianleach.co.uk" in a smaller black font underneath.