$7.00 U.S.

INTERNATIONAL

Spectrum®

THE MULTIVALUE TECHNOLOGY MAGAZINE | MARCH/APRIL 2015

Are You a Technology Bigot?

Also in this Issue

- Rocket U2 talks about UniData 8.1
- OAUTH 2 Login with MultiValue BASIC
- Clif Notes New Blood – Part 2

intl-spectrum.com

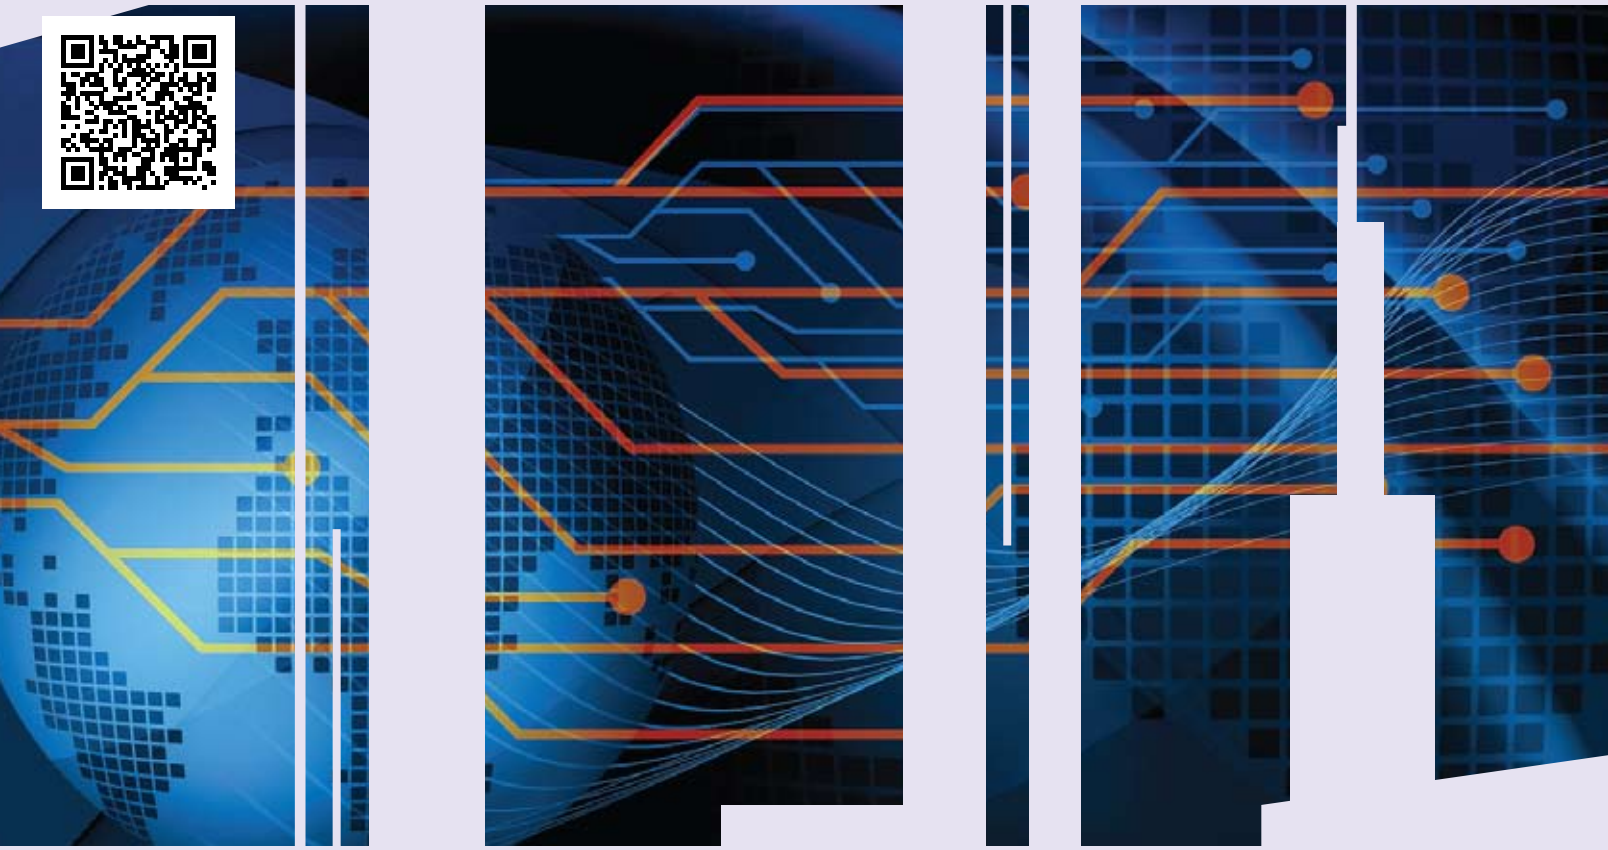

GET CONNECTED.

Knowledge and Education for the MultiValue Professional.

About Our Professional Membership

We are all busy in our day-to-day work and staying up-to-date with the current MultiValue technologies can be difficult.

Professional Memberships provide you access to knowledge, solutions, information, and code that you won't find in other locations.

Professional Membership Includes:

- Magazine in Electronic and Print Formats
- Newsletter
- On-Demand Videos
- Live Webinars
- Discounted Conference Rates
- Research papers
- Case Studies
- Source Code

INTERNATIONAL Spectrum®

http://www.intl-spectrum.com/membership/

INTERNATIONAL Spectrum
THE MULTIVALUE TECHNOLOGY MAGAZINE

6 **Are you a Technology Bigot?** When all you have is a hammer... Are we selecting technology for projects based upon what we like to use, or are we making objective choices? Kevin King attempts to resolve the age old question: Which technology is best? **BY KEVIN KING**

FEATURES | MARCH/APRIL 2015

9 **Rocket U2 talks about UniData 8.1** Rocket U2 has released UniData 8.1 with many new features and enhanced functionality. This release provides features such as true 64-bit Architecture, IPv6 integration, and Python integration. Charles interviews Vinnie Smith, Michael Byrne, John Zagnoli, Heather Smiles, and John Nunziato at Rocket Software to see how the release of UniData 8.1 provides enhanced value to UniData customers. **BY CHARLES BAROUCH**

14 **OAUTH 2 Login with MultiValue BASIC Overview** Consuming web services like Avalara Saletax within your MultiValue applications is easy, but many web services are starting to require OAUTH 2 login authentication. OAUTH 2 adds additional complexity, but implementation is much simpler than it seems when you read the documentation. This article will cover what is required to create an OAUTH 2 login process within your MultiValue applications. **BY NATHAN RECTOR**

16 **Clif Notes: New Blood – Part2** In part one, Clif discussed some of what that could be done to attract new blood to the MultiValue Community. In part two, Clif discusses not only how to attach Millennials to MultiValue shops, but also how to keep them engaged and interested. Millennials look at IT and the Enterprise from a totally different perspective than previous generations. With that in mind, read about how Clif suggests changes to the IT departments to address these conflicting perspectives. **BY CLIFTON OLIVER**

DEPARTMENTS

From the Inside page 4
From the Press Room page 12

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

NATHAN RECTOR
President
International Spectrum
nathan@intl-spectrum.com

The future is catching up with the power of MultiValue Applications, but are your MultiValue Applications catching up to the future?

We all know the power of the MultiValue data model. We also know the power of the Enhanced Stored Procedure abilities included with MultiValue Databases: BASIC, I-Types, and Dictionary Correlatives. But have you compared them to what your company is expecting of their datastore?

Challenge: Massive Data

What is *Massive Data*? Is it the size of the overall database or the number of transactions in any one file? Or is it the number of transactions consumed in a minute, hour, or day?

Challenge: Enhanced Decision Reporting and Alerts

Originally, databases were set up to store and retrieve data; nothing more, nothing less. Most applications are still designed to use the database for that purpose. But the emerging trend in database architecture is to start providing more advanced features in the stored procedure languages to provide answers to complex questions.

In the past, this approach was only available to the programmer who created the client application. And it required the client application to access massive amount of data, causing the process to be slow.

Nowadays, databases are beginning to take on even more complex logic. Oracle 10i touts the ability to write database procedures in Java. Microsoft SQL Server and PostgreSQL both support the use of Perl, Python, and PL/R for writing database functions.

Challenge: Hybrid Relationships

Hybrid Relationships are a method for correlating structured data with unstructured data. Not all data *can* be placed into well-defined structures. This is either because the structure is constantly changing, or because making it structured would require too many files, tables, accounts, and relationships to make it feasible.

Challenge: Data Any Place

Most people would define this as distributed data. People want to access the same data from any place, anywhere, using any device they have in front of them. The current buzz-term for this is "Cloud Computing" or "Cloud Databases".

While each of these technologies are implemented to address a specific issue, the real issues they are trying to address are "Ease of Access" and "Deployment of Changes."

Here is what the MultiValue databases are providing to address these needs.

Challenge Accepted: Massive Data

Thanks to the built-in compressed data format (Dynamic Arrays), MultiValue developers have been handling massive data for many years. Since we could always store large amounts of data and transactions in less space on the physical disk, we have never had to address the processing power and cost of storage issues that competing databases contend with. We're good with *Massive Data* as-is.

Challenge Accepted: Enhanced Decision Report and Alerts

We have always been sever-centric developers. Because of this, most of our programming, processing, and reporting have been done on the database servers, instead of offloading the data to the client.

MultiValue BASIC, I-Types, and Dictionary Correlatives provided us the ability to generate reports and analytics that the "traditional" databases are just now starting to implement.

INTERNATIONAL
Spectrum
MARCH/APRIL 2015

NATHAN RECTOR
President

CHARLES BAROUCH
Editor

SYDNEY BAROUCH
Editor

TRACEY RECTOR
Layout

Learn more about the MultiValue Symbol and see what MulitValue 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 Specturm, Inc., 3691 E. 102nd Ct., Thornton, CO 80229; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com. Copyright 2015 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. 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.

We have the tools to address *Enhanced Decision Report* challenges.

Challenge Accepted: Hybrid Relationships

This is not new to us. We have been creating data models which interact with structured and non-structured data for years. Since MultiValue files don't require a strict schema or dictionary in order to store data, we are not limited to structured data models. If we want to implement a partial structure within a non-structured file, we are able to do so using Virtual Dictionaries (I-Types and Correlatives).

Additionally, we don't have to store our data as two-dimensional data. We can create complex relationships without the complexity of a multiple file structure. MultiValue has always supported *Hybrid Relationships*.

Challenge Accepted: Data Any Place

While there is an ever-growing list of "Any Places," the need to create a *Data Any Place* access model is not new. MultiValue developers have been designing and creating "Cloud" applications and databases for years.

While the receiving applications still need to be built, attaching them is often as simple as building a web service. We have those tools on nearly every MultiValue platform; either built-in or with the help of readily-available third party tools.

So, next time you have a conversation with management about what a MultiValue database is or is not capable of, take a moment to think about what I've outlined here and say "Challenge Accepted." IS

AccuTerm™ by Zumasys

Please join Zumasys in welcoming AccuTerm to our growing family of MultiValue software products.

AccuTerm™ Version 7.1c AVAILABLE NOW!

www.zumasys.com/accuterm
818-951-1891

zumasys®
CONQUER THE CLOUD™

Open QM

taking multivalue ...
where it has never been before

- **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**

Ladybridge Systems Ltd
17b Coldstream Lane, Hardingstone, Northampton, NN4 6DB, England

Worldwide distributor: Zumasys, 9245 Reasearch Drive, Irvine CA 92618, USA
www.zumasys.com

www.openqm.com

Are You a Technology Bigot?

BY KEVIN KING

A couple of days ago, I was speaking with a friend who spent a good portion of the conversation trying to convince me that I should become proficient with a particular computer language. It would seem, based on this discussion, that this particular language is FAR superior to any other language that has ever been invented. It has all the features of this other language, and only the GOOD features of this other language, and not the baggage of THAT OTHER language, and, and, and it does all this neat stuff, and… It's just so cool I simply wouldn't believe it.

I have to admit this conversation made me chuckle a bit. While the information was certainly useful and the banter about "which peanut butter tastes more like peanuts" kind of discussion is always entertaining to some degree, what he didn't know was that I had spent a greater part of the day in similar discussions about other languages and platforms. Which database is better? Which language is better? Which operating system is better? Which development environment is better? Which technique is better? Considering that a good deal of time seems to be spent going back-and-forth over these very issues, I've decided to take a departure from my favorite topic (that is, our Aspen database project) and attempt to answer these questions once and for all.

…the answer is generally hidden in a wealth of subject-matter knowledge waiting to be mined through intent, focused listening.

Sounds pretentious, doesn't it? How can one guy purport to provide an answer to these age-old debates? Forgive me if I offend any delicate sensibilities, but the answer has always been right under our noses...

Listen.

What's the best database for any given project? Listen to the person who will use the information. What's the best language? Again, listen to the person who has the problem to be solved. With either of the preceding questions, and the numerous questions that haven't been specifically mentioned, the answer is generally hidden in a wealth of subject-matter knowledge waiting to be mined through intent, focused listening.

"Hey now," I can hear you say, "my customer knows nothing about languages or IDEs or platforms or operating systems. How can listening to my customer help me select the best language for a project?"

While this may very well be true, the fact remains that your customer has a

certain level of subject-matter expertise. He or she holds the keys to the problem that they want you to solve. By understanding their specific problem from their perspective, and comparing this understanding to your understanding of various technical options, you arrive at a better position to select the most appropriate tools for the task.

Unfortunately, this kind of listening is becoming more and more rare. Instead, the technology community at large seems to be mired in a form of "technology bigotry," where the only solution to a given problem is whatever technology happens to be en vogue at the time. Perhaps you've heard comments like these? (Substitute any technology in place of the blank, i.e. Java, Unix, or even Pick.)

- "We're a _____ shop."
- "_____ has a perfect solution to any problem."
- "_____ is the only REAL (database | language | operating system | et al)."

and my personal favorite:

- "Real programmers never use _____."

This kind of thinking tends to fall under the category of the old cliché: When all you have is a hammer, everything looks like a nail. Problem is, not everything is a nail. Occasionally we need wrenches, sockets, and screwdrivers to get the job done. The same holds true in the software world. Everything doesn't always fit into a nice little box that can be solved by one and only one option; sometimes a task calls for different languages, databases, environments, and maybe even a dip into a new operating system. The trick, if there is one, is simply to have enough tools to make an informed choice possible.

This is where it can get dicey. While I can't speak for the community at large, I suspect that most people in the business of creating software are pretty busy — busy enough to where there's not a lot of time available for exploring new options.

Then there's the whole cost factor. Learning new options requires some kind of investment; whether time, money, or both. And: Is there really a benefit to knowing the latest whizbang technology when you may or may not actually use it?

In answer to that question, it has been my experience that it is ABSOLUTELY beneficial to continue to learn new things, even if you may not ever use them. Exploring new options almost always provides new ways of looking at problems and situations. Who, after learning something new, hasn't looked at things just a little a bit differently? It is these perspective shifts — some moderate, some dramatic — that provide the foundation for making better technology choices in the future.

Sometimes when we learn new things, it provides evidence to support what you may already believe. How many of us (and you don't have to raise your hand) have made a decision against a particular technology or product simply because we heard from somewhere that it had this problem or that problem? Did you personally experience that problem? Did your source happen to tell you that, while this or that problem may exist, there was this or that way around the problem? A solution that was better for a variety of reasons?

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@eslny.com

Consultants Needed Nationwide

No, of course not. Technology bigotry works like that; in each generation the prejudice grows stronger and stronger while the evidence to support it becomes less and less compelling to a point where fact and fiction become indistinguishable. At this point, the solution providers — that is, each of us — are on perilous ground. The discounted technology may be just what the customer needs, but we're too blinded by prejudice to see it for what it could be in that context. Perhaps your original technology choice was the best one, but how can you know for certain unless you've personally experienced other options?

In closing, my point is this: simply try to interject some reason into the ongoing debates of "best" anything. Defining "best" without defining a context is like saying yellow is the best color. Perhaps for some uses it is. For others, it's not. The same is true for technology; there can never be one best anything without considering the context. Listen to your customer. Have a well-equipped arsenal of options from which to choose. Use all this information to make the best choice.

After all, the software business isn't really about creating software; it's about solving problems. IS

KEVIN KING is the President and Chief Technologist with Precision Solutions, Inc., a leading technology solutions provider in Longmont, Colorado. He can be reached by email at Kevin@PrecisOnline.com or by voice at 303/651-7050.

mv QB

QuickBooks API for the MultiValue Database

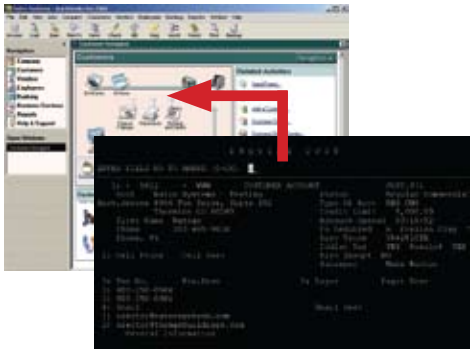

- **Read/Write Directly to Quickbooks Databases**
  Customer, Vendor, Invoices, Purchase Orders, Chart of Accounts
- **mvQB API is Designed for the MultiValue Program to Use**
  All routines are simple BASIC calls designed for the developer. No special user interfaces required.
- **No Need to Learn the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

ROCKET U2 TALKS ABOUT UNIDATA 8.1

BY CHARLES BAROUCH

We sat down with Vinnie Smith, Managing Director R&D, Rocket U2, John Nunziato, Product Manager, Rocket MultiValue Databases, Michael Byrne, Rocket MultiValue Product Evangelist, John Zagnoli, Rocket U2 Development and Engineering Manager, and Heather Smiles, Product Marketing Manager, DBMS and Application Servers, to talk about UniData.

IS: ***Spreadsheets and Word Processors reached a point where the supposed added benefit of each release was less and less useful. Tell us some powerful, practical things that companies will get out of the new UniData release.***

VINNIE: 8.1 is the biggest release in a decade. There's a number of key areas of benefit. The 64 bit architecture is more stable.

JOHN N: The obvious effect is on the file access size; being able to address much larger files. We can now exceed the two gigabyte boundary. That's brought along reliability improvements. The process architecture now can access a 64 bit memory model. You can process much larger datasets in memory.

MICHAEL: Account-based licensing allows administrators to reserve a certain amount of licenses for each virtual group. That way you always have enough licenses for a specific group, like admins. This really crosses over into our hosted solutions, especially anyone who is doing SaaS (Software as a Service). And it's in a configuration file. You can configure accounts into virtual groups. Everything not in the file is in the default pool. If you sell cloud services, you can control the split of the licenses between clients.

VINNIE: This will help in building a multi-tenant solution.

JOHN N: Security. SHA1 and SHA2 security on files. We also have improvements to protect against Heartbleed and Poodle attacks.

JOHN Z: Automatic data encryption at rest includes the keys and indexes. You can 100% encrypt everything in the data file. If you store data in your dictionary, you can encrypt that, too.

MICHAEL: That protects you in the case of your disk drives being stolen, or your data backups.

VINNIE: And supporting IPv6. Because, as the world runs out of Internet addresses...

IS: ***We hear "IPv6" a lot, but I'm not sure everyone understands what it is. Can you explain it briefly?***

JOHN N: IPv6 uses 128 bit, that gives us 3.4x1038 IP addresses. [Editor's note: IPv4, by comparison, tops out at roughly 10 digits (decimal) worth of addresses vs. IPv6's 32 digits (hexadecimal) worth of addresses.] More security, including the authentication header protocol and the simplified packet header structure, which makes data packets faster and more efficient. More data security — encryption capabilities — when communicating between systems. We now offer two BASIC commands: SET-IPV and GET-IPV to enable and disable IPv6. Or allow both IPv4 and IPv6, prioritizing them. There will also be config file options.

VINNIE: Basically a dual stack capability incorporating both protocols. Calling out to other IPv6 devices — the Internet of Things — being able to call out to them is very important in terms of extending what you can do with UniData.

IS: ***The big fear in any upgrade is stability. Can you tell us a little about your testing process?***

JOHN Z: This was a major code upgrade, especially the 64 bit extensions. We had cross-function testing, QA testing, Support Department testing, and beta testing. Beta was August and ran through to January. After that, we had several customers run their existing applications and found no regression errors. We got acceptance from everyone in the beta.

IS: ***How does the upgrade work? Has the process been improved?***

VINNIE: Especially in multi-system sites, upgrades can be quite challenging. We've automated out most of the manual steps. You can even add your own post-upgrade scripting. We've done our best to make the upgrades easier for everyone. We are quite happy that we've smoothed the upgrades for everybody.

IS: ***I know that when I work on a large, complex software project, some aspect of it becomes a pet project, something that I'm particularly proud of at the end. Do any of you have a pet feature in UniData 8.1?***

JOHN Z: I can name two. I'll start with the local subroutine. If you have a large mainline program and you want to add a subroutine, you can now build it into the body of the program using variable scoping. Everything compiles as a single unit. We had already released it in UniVerse 11.2. We extended it into UniData as well. It's an internal subroutine but it has separate variables like an external subroutine.

MICHAEL: It helps with repetitive work where you might otherwise risk stepping on global variables.

JOHN Z: The other is the replication performance monitor. You can see each file and how it is changing: data utilization, buffering... This is at a level you haven't been able to see before in any other monitoring tools. (See video interview with Jonathan Smith at: https://www.youtube.com/watch?v=t6NzEd2XXio)

MICHAEL: Speaking of pet projects, Python isn't in this release, but 8.1 gives us the groundwork. We are laying the foundation for Python while thinking about how we can extend that architecture into other languages. I find that exciting to me; getting connected to the outside world. Making our databases available to a whole new audience of developers who know other languages.

IS: ***UniData, UniVerse, and D3 are all part of the Rocket family now.***

Is there a feature or aspect that sets UniData apart?

MICHAEL: In my opinion, from a reliability standpoint, UniData has always had RFS and Replication in the engine. We have several financial — big customers — that depend on that.

The database servers are all somewhat similar, especially as we cross-implement features to make them all more robust. We're really thinking forward, as we develop for one engine... thinking of how it will be implemented in the other databases.

HEATHER: We also have mvBase, too.

VINNIE: We think in themes, like security, compliance, reliability. We try to love all our children equally, so to speak. A lot of our code is moving to a common code base in the middle tier — client connect — aspect. On the Rocket website, there are white papers about 64 bit and other features. We also have lots of examples on Github. You can also follow us on LinkedIn, Twitter, and Facebook.

IS: ***Anything "deeper-dive tech" that you wanted to talk about in this interview?***

JOHN Z: Dynamic files. Not only are our dynamic files 64 bit, but we've added a new hashing algorithm. We've also improved the merge and split process. It has been tested against extremely large files.

JOHN N: For replication, we added data compression — LZ4 — to reduce the footprint between publisher and subscriber by at least 50%.

IS: ***Are there any other new UniData features that you'd like to talk about?***

VINNIE: We covered the high-level list. I do want to point out that we have an upcoming MV University. We'll be discussing many of the 8.1 features in more detail.

HEATHER: We've already sent out a save the date for Harrah's in Vegas. 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, PerehelionSF, and the Interrogative series, which begins with Tiago and the Masterless.

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 • http://sjplus.com

prc®

WANT TO SEE A SPECIFIC TOPIC?

InternationalSpectrumislooking 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

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

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

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

From the Press Room

Vermont Teddy Bear Relies on Rocket U2 to Improve Service and Increase Revenue

The Vermont Teddy Bear Company is the largest handcrafter of Teddy Bears in North America. Each BearGram® gift includes a customized Vermont Teddy Bear accompanied by a personalized greeting card and free candy treat, all packaged in a colorful gift box with an air hole.

The Vermont Teddy Bear Company (VTBC) began nearly 30 years ago offering a single product line that included a teddy bear, a custom outfit "kit," candy and a greeting card. VTBC has grown dramatically, reaching a peak of over 47,000 orders and more than 60,000 packages shipped in a single day. Staffing fluctuates along with sales, expanding from 150-180 year-round employees to as many as 1,000 during peak months of December, February and May. The continually expanding product line and combined volume across these various divisions have helped the company achieve growth of over $50 million in sales.

"We couldn't do what we do without this MultiValue product. It's great for developers. There are a lot of nicely written pieces I can use without having to create new code." — Bob Stetzel, Vice President of Information and Ecommerce Technology

Because its sales are closely tied to gifting holidays, VTBC experiences peak periods when orders grow from several hundred per day to many thousands per day (e.g. February 13th, the last day to order a gift that arrives on Valentine's Day). In the early days, they were using an outbound telemarketing system that had numerous issues, and knew they needed a better solution to handle huge order volumes in the face of deadlines, get them out on a timely basis, and manage customer service and shipping.

Get full case study at:

http://www.rocketsoftware.com/sites/default/files/resource_files/Vermont-Teddy_cs.pdf ■

jBASE 5.2.30 Released

Zumasys is pleased to announce their first software release since combining jBASE Development and Support. At jBASE, they are committed to the continuous improvement of their products and support services by implementing recommendations and enhancements that are suggested by users and that align with industry standards and trends.

The following fixes and improvements were made in release 5.2.30, based on feedback and feature requests from you:

- jBASE license validation utility
- Enhancement of the PH-KILL command to set the End Date and End Time
- Supported SQL command documentation for jODBC
- Enhanced jODBC connectivity ■

Ashwood Computer Inc. Is offering a NEW Python to jBASE Connector

Programmers can now write web applications using Python and then directly read and write jBASE files from those applications.

One customer, doing medical case management, is using the new connector to put a browser front end on a portion of their internal application. They are also using it to integrate their jBASE application with the company's web portal and some of the Google / Chrome applications they use. Users may provide their own Python/Django developer talent, or contract directly through Ashwood for the services they desire.

This new connector is offered with several different levels of consulting services. You can choose from basic setup and installation services, hourly services and assistance as desired, or contract for a complete customized implementation and programming services offering. ■

QM Release Notification (3.4-1)

This is a minor release containing the following new functionalities requested by our users...

Logging in a telnet or QMClient session on Windows now supports domain\username as an alternative to username@domain .

The initial state of terminal input case inversion can now be set using the INVCASE configuration parameter.

The CREATE.ACCOUNT command now has a NO.CASE option to create the VOC as

a case insensitive file.

The GET.LIST and SAVE.LIST commands use a default name if no list name is given.

The LIST.PRINTERS command can be used to show a list of all printers known to the underlying operating system (currently Windows only).

The CLEAR.SELECT mode of the OPTION command causes the default select list (list 0) to be cleared on return to the command processor except when the list was created by a command or program whose role is to create a list for processing by subsequent commands.

The NUMERIC.EXPONENT mode of the OPTION command enables acceptance of exponent form numeric constants (123.45E-1). This applies whenever QM converts a character string to a numeric value.

The $QUERY.DEFAULTS record has been extended to allow specification of the format code to be applied to the automatically inserted @ID column if there is no dictionary entry.

The $BASIC.OPTIONS record now supports turning off compiler modes. This is useful if a shortcut name is used that sets many modes but some are not wanted.

Include records to be inserted automatically into all programs can now be specified in the $BASIC.OPTIONS record.

The QMBasic INPUTNULL and INPUTTRAP statements have been added.

The QMClient API QMIConv() and QMOConv() functions now have multivalued equivalents, QMIConvs() and QMOConvs().

The QMClient API QMIndices() function can be used to retrieve information about alternate key indices. ■

Rocket Software is proud to announce Rocket UniData 8.1

64-Bit File Support - Experience increased I/O performance and work with larger file sizes than ever before! No more 2GB limit on file sizes!

64-bit Architecture - Take advantage of the most current hardware without the memory limitations posed by 32-bit software! Eliminate the need to maintain a mixed-architecture environment and see a potential performance benefit when the database is fully cached in memory.

Support for IPv6 - Be ready for the Internet of Things! As the world is running out of IPv4 Internet addresses for all web-enabled devices, IPv6 compatibility provides the ability to create more IP addresses for years to come, and utilizes a more secure data pack in the process.

Replication Performance Monitor - Allows system administrators to monitor traffic, system state, and metrics across multiple publishers and subscribers. Identify performance bottlenecks in real time that could be tuned by changes made in the repconfig settings.

Account Based Licensing - Allows application service providers to manage licenses and allocate based on the account being used. This way you can prevent lower priority accounts from consuming licenses intended for higher priority accounts.

Install / Upgrade Improvements - UniData 8.1 has overhauled its installer and update process in this release, with significant customer input into the design changes. See the UniData Installation Guide for full details!

Security Enhancements - Secure your data with the latest in OpenSSL updates, @ID & Index encryption, and other Automatic Data Encryption improvements. ■

OAUTH 2 Login with MultiValue BASIC Overview

BY NATHAN RECTOR

Many web services are starting to require developers to use the OAUTH 2 authentication flow in order to access their web APIs. While this makes working with them a little more difficult, it isn't as bad as some of the documentation makes it out to be.

If you are wondering why companies are doing this, it's pretty simple. Companies like Google, Facebook, Linkedin, Salesforce, and Microsoft have all spearheaded the use of OAUTH 2, causing many other companies to follow close behind.

OAUTH 2 replaces the original OAUTH 1 URL signature process. This is good since OAUTH 1 ended up being a more complicated signature to implement. Some companies, such as Twitter, still use OAUTH 1, but even they have been slowly moving over to the new OAUTH 2 standards.

OAUTH 2 requires developers to include additional information with every web service call to verify that the person making the request is truly who they say they are. This additional information allows the API to determine if the person and/or device that is issuing the call is allowed to do so.

OAuth 2 requires developers to include additional information with every web service call, verifying that the request is coming from an authenticated process or person.

Since OAUTH 2 protocols can distinguish the device in addition to the user, the API can de-authorize a device without de-authorizing the user. As you can see, this has many security advantages.

Login Flows

There are several login flows available with OAUTH 2. This multiplicity is one of the reasons the standard seems so complicated. Another reason is that most of the documentation covers access via websites and mobile devices, which introduces much of the confusion.

All login flows require that the API provides a *Consumer Key* and a *Consumer Secret.* These values must be stored within your application and are used to identify the process or device. These keys are usually assigned after you fill out an application to use the web service.

Once you have the *Consumer Key* and *Consumer Secret*, you need to find out which login flow the web service is requiring you to use.

These are the three most common login flows:

Password – Resource Owner Password

This login flow is not as secure as the next login flow, but it is the easiest to implement. You provide the *Consumer Key* and *Consumer Secret* with a username and password in order to get an access token to use later.

Native – Desktop/Mobile Device Authentication

This login flow requires you to have access to a webserver. During the login process you must include a URL, telling the API where to send the access token. This means there is no user

interaction or data entry, but it does require your LOB application to have its own web service available to receive the access token.

As always, there is an exception to this. Some web services allow you to supply a special callback code called an *oob*. This tells the web service to redirect to a page that includes a PIN number. A user must enter this PIN number.

A second authentication is done using the PIN number to verify that everything is okay. This requires that API developers provide this functionality, so sometimes you see it and sometimes you don't.

Limited – Limited User Access

This login flow allows the authentication process to happen without a callback URL found in the "native" login flow. The application uses a "polling" function to request the access token.

For LOB applications, this would be the next easiest to implement. The only problem is that the developers of the APIs must allow this type of login flow, and most do not.

Access Tokens

The whole point of the OAUTH 2 process is to issue a temporary access token to a device/application combination. This access token can then be de-authorized by the API at any time, or the API can force the device or application to re-authenticate after a period of time.

Along with the access token, the API usually provides an *expire-in* value; usually provided as the number of seconds for which the token will be accepted. To figure the exact date and time when the token will expire, your application needs to record the date and time when it received the information. Otherwise it can't check to see if the current date and time are larger than the expire-in value.

Most APIs include a refresh token as well. Refresh tokens allow you to speed up the re-authentication process. There is a limited amount of time during which a refresh token is valid. This will vary from API to API, so it is important to be aware of this, especially if the API is not used often. If the refresh token has expired, you will have go through the complete login flow from the beginning without the benefit of the refresh token.

Once you have the access token, it will need to be added to every request to the API from that point forward.

URL MetaData

Before issuing the request to the API, you must add a new header field to the URL information. This field is an 'authentication' header, and looks something like this:

```
Authorization: bearer aopi991/
all93k1jf7
```

This is the information that the API will use to validate that your application is allowed to access this specific service.

Next...

My next article will talk about the control record storage and BASIC code needed to comply with OAUTH 2 authentication. IS

NATHAN RECTOR
President
International Spectrum
nathan@intl-spectrum.com

New Blood — Part 2

BY CLIF OLIVER

In part one, we discussed some of the things we could do (and a number of things not to do) to attract new blood to the MultiValue community. Now let's talk about some of the things we need to do, not only to attract Millennials to our shops, but also to keep them. It does us little good as a community if a bright, energetic, young developer joins our team, suffers their own version of "burnout," and two months later quits because they are bored to tears. Or if they are embarrassed to admit to their cohorts what they do for a living.

The first thing we need to understand is that Millennials are coming from a different perspective than a lot of the previous generation did. (And definitely different from the previous-previous generation.) For example, a job is not just a job. It is not just a way to make money in order to buy the Beamer or the newest catamaran. Millennials expect jobs to have "meaning." Now the individual's sense of what "meaning" means is as varied as it is in the previous generations. But whatever it is to a particular individual, as a group they expect a job to be more personally fulfilling than simply collecting a paycheck and busting their butts for a promotion. Remember, Google is your friend. Go out and do some searches and you'll find numerous articles about how this group is frequently willing to take a lower paying job that has a better working environment, better social interaction, and what they perceive as a higher company social conscience.

Now unless you are in an upper management position in your company, possibly even a C-level, there's not a whole lot you can do about the company's social conscience thing. But there are some things middle management can affect that will determine whether or not a newly hired millennial continues as a MultiValue developer. Everybody on the development team plays a part in determining whether the new developer will stay or bolt.

Let's take a look at what the younger crowd thinks about another "big mainstream" environment as to its desirability for even a short term employment, let alone a long-term career.

You can see it here, in a couple of posts that I found right off the bat while I was doing web searches. (In other words, they popped up in the first two screens. I didn't have to go digging all night for these.) They show why young people don't want to work with mainframes. As you read this I am going to ask you to mentally substitute the word MultiValue for the word mainframe.

All the italic emphasis is mine to highlight the points I think should slap MultiValue developers in the face. A lot of what we are accused of — old technology, outdated programming tools, dinosaur mindsets — are thrown at the main-framers, also. So take a look at these two posts about why young developers want to avoid mainframe jobs.

> "In a competitive IT world if you spent the couple of years it would take to actually come up to the beginnings of reasonable speed in mainframes, you wouldn't lose your previous skills but *you'd automatically be less appealing when looking for other work*, just like if you'd spent two years doing forestry or managing a Starbucks: *looking like you're out of the loop even a little bit does you no favors when being compared to someone who doesn't look that way.*"
>
> – Matthew Frederick May 12 '11 at 9:36
>
> "I'm a young programmer. I've never seen a mainframe, never had a sandbox/virtual mainframe to play with, *never had a friend come up to me and say, "This is really cool, check it out!"*. I see the web every day, there's readily available - and free - webapp dev learning tools, and *all my friends are doing neat stuff in it. Which am I going to choose?*"
>
> – Beekguk May 11 '11 at 19:51

Do you see what I mean? How on earth do you expect to retain a new

developer who has had a computer and always-connected Internet most of their life? A developer who spent their time in school studying to work in that modern environment of point and click, drag and drop, always available websites, and now mobile apps? Do you really think that they are going to be satisfied working in a shop day in and day out coding in BASIC while their friends are enjoying constantly learning about new technologies and working with the latest methodologies and frameworks?

Newsflash, people. You can't pay them enough to do that.

Fortunately, we are now at a state of affairs on most of the MultiValue platforms that we don't have to. There is no reason on earth for a shop to keep developing telnet-based green screen programs when there are a number of modern tools available for desktop, web, and mobile applications. But yet I continue to hear the same old arguments about modernizing things being too expensive, less efficient, etc. that I have been hearing for more than the last 10 years. (Yes, I know there are exceptions – different discussion.)

You expect a 25-year-old whiz-bang developer to stick around in that type of an environment?

Now I will admit that we have some things going against us in this battle for young minds that we really can't do anything about. For example, the very name of our programming language. BASIC. It doesn't matter whether we call it MultiValue Basic, Data/Basic, U2 Basic, QM Basic, or whatever. That B-word will kill us every time. As Brian Leach pointed out in one of his recent Spectrum Magazine articles, even Microsoft's Visual Basic has the same baggage of being considered a toy language because of the B-word. "Serious".NET developers work in C# and leave Visual Basic to the hobbyists.

Well, I'm afraid that's just something that we're going to have to live with for now. There have been a number of suggestions over the years of how to deal with that. But when you consider the cost and resistance of changing a name that has its roots in the 1970s, well, good luck with that.

But one of the things that we can do is to present it to the younger developers as not being the primary development language for everything but more like a server-side scripting language. There is no reason why the new developers can't be developing in something else (C#, Python, Java, etc.) on the client side while the traditional MultiValue developers are providing them with subroutines on the server-side to call for the data access and business logic.

Now we certainly hope that we can entice at least a few of the young developers into learning the server-side language. We need them not only to be able to do their own server-side programming when people retire or move on but to help maintain some of the "inherited" code while we are slowly modernizing everything. So by whatever means, we need to entice them into learning this programming language. So we sit down with them, edit a program, and start to go through it, explaining the features of the language.

After recovering their breath, picking their jaw up off the floor, they are likely to go running for the HR department to turn in their resignation. Yelling as they run down the hallway, "are you <explicative deleted> kidding me!? THIS PRODUCT HAS ITS PROGRAMMING DONE IN ALL UPPERCASE!".

Well, you've just lost another one. All of your talk about how advanced and easy the programming language is to use has been blown out of the water and forgotten because it looks at first glance just like COBOL, FORTRAN, BASIC, etc. In other words, a language that will destroy your career and make you unemployable if you work in it. Just perception? Maybe, maybe not. But to most people, not just Millennial new software developers, perception is everything.

How many times have you heard someone in the sales department quote the old line, "You sell the sizzle, not the steak?"

We will talk about this issue in part three. IS

International Spectrum Magazine

on Paper, Online, and Beyond

International Spectrum Magazine is moving beyond print and is now accessible on your favorite e-readers and mobile devices. By providing International Spectrum Magazine in the popular e-reader formats, subscribers are now able to take the magazine with them everywhere they go.

Let us help you keep your MultiValue skills up-to-date and efficient.

Now on Kindle, Kindle Fire, Nook, e-Readers, iPad, iPhone,and Android Tablets!

INTERNATIONAL Spectrum

http://www.intl-spectrum.com/magazine/