

\$7.00 U.S.

INTERNATIONAL Spectrum[®]

THE MULTIVALUE  TECHNOLOGY MAGAZINE | JULY/AUGUST 2015

CRAZY PEOPLE

Everyone is crazy but you!

ALSO IN THIS ISSUE

- Business Intelligence in the MultiValue Community
- Roll Your Own (Highly Experimental) MultiValue Database — Part 2
- Rolling Your Own MultiValue Web Connector - Part 1

April

11 - 14

2016

The Premier MultiValue User Event



From application development and technical systems support to complex business and web development, we offer complete IT solutions.

RETURNING TO THE WIGWAM , LITCHFIELD PARK, ARIZONA



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.

Topics covered at the International Spectrum MultiValue Conference include:

- **Integration**
- **Database Management**
- **Mobile Integration**
- **Reporting and Data Presentations**
- **Project Management**

...and more.

REGISTER AT

www.intl-spectrum.com/conference





6 Business Tech: Crazy People The world demands synergy between business and technology, but the conversations between the two sides are often shouted across a gulf chocked full of misunderstanding. In the process, both parties sound crazy to each other until a commonality can be found. Chuck talks about how to address these issues and suggestions on how to find the commonalities so you no longer sound like “All those Crazy People”. **BY CHARLES BAROUCH**

9 Business Intelligence in the MultiValue Community Business Intelligence is certainly not a new concept. It has been utilized by businesses for almost two decades. Due to the lack of native support for the non-first normal form database structures, it has not been an easy journey for the MultiValue community to adopt most Business Intelligence (BI) solutions. As a result, some MultiValue application providers have taken the strategic approach of using various ETL (extract, transact and load) processes to migrate the MV application data into a relational data warehouse or data mart. The end user community then performs their analysis on this derived data source. Dorien Gardner discusses other solutions and means to access MultiValue Business Intelligence without remapping into non-MultiValue database. **BY DORIEN GARDNER, ROCKET SALES ENGINEER FOR ROCKET D3 DBMS AND MVBASE**

15 Roll Your Own (Highly Experimental) MultiValue Database — Part 2 In Part 1 of this series we laid the groundwork. In Part 2 we will take a look at some basic file transaction functions. A database isn't a database if you can't CRUD (Create, Read, Update, Delete) it up. **BY CHARLES BAROUCH**

18 Rolling Your Own MultiValue Web Connector - Part 1 In case you have been asleep the last few years, let us fill you in on some important news. The Web has taken over. In order to meet your customers' demands and stay competitive, the data in your MultiValuebased applications must be accessible by Web technologies. While there are a number of fine products in the MultiValue world to allow you to do this, if you are willing to get your hands dirty with a little bit of programming, you can accomplish a lot of this with tools readily available to you as open source. **BY KEVIN KING**

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

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 2016 Spectrum, it will be April 11th-14th in The Wigwam in Litchfield Park, (Phoenix), Arizona. We are having it at the same venue in which the 2014 conference was held. While it's not the PGA, it also has a beautiful golf course, so bring your clubs!

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 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, Email, Spreadsheets... there are versions that are free and versions for a fee. All can be

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. 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.



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



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 been an island unto ourselves for so long, we sometimes forget that our business has information that other systems can use, or are 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 trends and topics that you would like to have covered in sessions at Spectrum? I would like to hear from you.

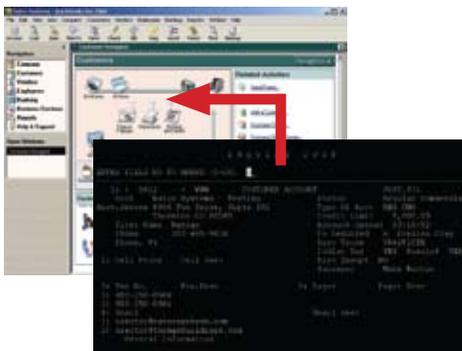
Email me: nathan@intl-spectrum.com



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



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**



Providing Solutions to your MultiValue Questions

Phone: 303.465.9616
E-mail: mvqb@natecsystems.com
Website: www.natecsystems.com

Crazy People

BY CHARLES BAROUCH

It's Just Solitaire

I've been conditioned to put the aces up in a certain order when I play solitaire. Hearts first, then spades, then diamonds, and finally clubs. Because of this, I get a small positive bump in my mood when each of the aces arrive in that order. The last time I played, the heart ace, spade ace, and diamond ace each came up the "right" way. What's the chance that the next ace will be a club?

Some of you re-read that last sentence to be sure that I wrote something that stupid. Stay with me, there's a reason I phrased it that way. Another one: My budget for hardware was over by ten percent. If I spend another thousand, will that fix the problem?

Meme What You Say

There is an Internet meme, which I will restate in a family-friendly way, which says "The problem is that educated people sound like crazy people to uneducated people." What many of us don't realize is that the reverse is also

The world demands synergy between business and technology, but the conversations between the two sides are often shouted across a gulf chocked full of misunderstanding.

true. Uneducated people sound like crazy people to educated people.

Each of us has been in workplace conversations where people have asked us questions that sound a lot like "What's the chance that the next ace will be a club?" Our instinct is to see these questions as crazy, and by extension, question the sanity of people who ask them. So, we sound crazy to them — because we're smart — and they sound crazy to us because they are dumb... except that's not exactly true.

We are the They

Everybody is uneducated on some subjects. Some of us can calculate

compound interest in our heads. Some of us can convert between hex, decimal, and binary while carrying on a conversation. Some of us can pick out individual notes while listening to a symphony. Don't get me wrong, there are some people who are irredeemably stupid. However, that's not the way to bet.

There are three glaring problems with the "they be stupid" approach. First, it's an unfixable, and therefore nearly unwinnable position. Second, if it isn't true (or even if it is) you are insulting someone whom you are trying to sway to a new perspective. Third, and most importantly, it might be us being stupid.

Here's the thing about that third option: I was on a sales call — I was the support person accompanying the salesperson — where the sale was in jeopardy because of the example used in our demo. The client knew that Sears would never order that style of socks. For her, the fact that the salesperson suggested that particular scenario left

her too gobsmacked to respond. His take: He's the educated person (knows our system) and she's the uneducated person looking at him as if her were crazy. Her take: He's the uneducated person (doesn't know the business) looking like a crazy person to her because she is educated on the customers' patterns of action.

Hide in Plain Sight

I've had the opportunity to sit in a lot of other people's office while business was being conducted around me. Being a consultant offers a certain level of invisibility which can be very educational. When I listen to the after-conversation — what happens after the customer hangs up the phone, or the prospect leaves the presentation — I almost always hear a complaint about stupidity.

As an outsider, I find it hard to keep from laughing. The gripe usually sounds something like this: "She thought you could combine those policies? Doesn't she know one is personal and the other is corporate?"

The funny thing is, if she's running a Sub-S or sole proprietorship, the line between the two is very blurry. This is an example of each side being uneducated on parts of the same conversation. You would be shocked at the number of times I've heard this sort of judgmental response. I'm sure if I followed the prospect out of the room I'd hear much the same from them.

Technology as a Foreign Language

This, by the way, is why my column is called Business Tech. The world demands synergy between business and technology, but the conversations be-

tween the two sides are often shouted across a gulf chocked full of misunderstanding. It would be nice if both sides became more educated, but there are only so many hours in a day and, hopefully, they are not all dedicated to work. We might not have the time or opportunity to learn their business down to the nuances. They probably won't be learning how to code and do analytics any time soon. When the chance to learn does present itself, we need to avail ourselves.

Much of my success in technology comes from the fact that I didn't set out to be a technologist. I set out to be either a writer/teacher or a business executive. I come to tech with a prejudice toward seeing the business case for everything I do. I come to it with a need to explain everything. Even when I don't know the specific business, I



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

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

still try to think from the business perspective.

Most of your bosses, odds are, have never known exactly what you do for a living. That's a staggering idea that we've become numb to over time. It sets a precedent for miscommunication because our jobs are largely opaque to the outsider. It is assumed that we are hard to understand because we do that — insert miracle here — which keeps things running.

Looking Back

Think about your own experiences and realize how many unpleasant situations in your career were caused by fundamental mismatches of expectations. Sally in legal knows that you can't charge a client for unshipped goods in a retail setting. You know that there's a time limit between authorization and settlement on a credit card. But, since you don't know the legal-

ity, and she doesn't know the technical constraints, things can get contentious. Bill walks into Alice's party planning office, looking to put together a massive family party. She's talking ham vs. bacon to Bill, who is unable to eat either for religious reasons. He tells her, so she assumes kosher — which she can accommodate through a partner — when the reality is that he is following the rules of halal, and area outside her expertise. On the other side, Bill keeps talking about amounts of chicken and beef which exceed the available refrigeration space, a fact that Alice has firmly in her mind because she knows her partner's equipment... and round it goes.

A Second Look

My budget for hardware was over by ten percent. If I spend another thousand, will that fix the problem? Guess what... the answer might be yes. Depending upon the size of the budget, and what level of spending it takes to trigger a volume discount, that additional thousand might tip things, resulting in a decrease in the total cost. Some questions are smarter than they seem. It comes down to education; try to be educated. Learn all you can about the business your tech supports. If we're going to be thought of as crazy, we want to be the smart kind. **IS**

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

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.



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.

Business Intelligence in the MultiValue Community

Business Intelligence is certainly not a new concept. It has been utilized by businesses for almost two decades. Due to the lack of native support for the non-first normal form database structures, it has not been an easy journey for the MultiValue community to adopt most Business Intelligence (BI) solutions. Almost all the top solutions require SQL access via ODBC or JDBC. Rocket® Universe, Rocket® UniData, and Rocket® D3 DBMS all support exposing the databases as SQL-compliant. However, this comes with a price in performance, maintenance (SQL Mapping), and knowledge of a language that is not natural for most MultiValue professionals.

As a result, some MultiValue application providers have taken the strategic approach of using various ETL (extract, transact and load) processes to migrate the MV application data into a relational data warehouse or data mart. The end user community then performs their analysis on this derived data source. This approach is certainly valid. There certainly are success stories from some of our partners. On the surface, this may seem to help mitigate

Transaction write times can be dramatically impacted if too many indexes are maintained.

the data access dilemma; however, as most companies soon realize, using this approach to Business Intelligence can be costly, as companies must:

- Purchase and maintain a relational database
- Add or internally develop expertise in the additional database
- Add or internally develop a strong knowledge of SQL
- Increase expenditures on dedicated hardware for the data warehouse or data mart
- Develop and maintain the ETL processes to normalize and load the relational database with the MV data
- Dedicate resources to a single purpose solution

Although this is certainly not an exhaustive list of what is needed for this deployment architecture, it is easy to ascertain that this approach adds quite

a bit of expense and complexity to the overall application solution. Another critical impact of this approach is that it naturally increases the likelihood that the data being analyzed may be too latent for time-critical decision making. This is directly impacted by how frequently the synchronization between the transactional/operational systems and the data warehouse or data mart occurs.

Native Access Data in a MultiValue Database

One of the ironies of making an MV database SQL-compliant, either within the MV database or migrating the data to a relational database, is that you essentially have to break apart the natural relationships of the MV database design by normalizing the nested structures, only to join them back together via SQL queries. This essentially requires that either the end user has knowledge of the table relationships or that highly skilled people have to develop a metadata layer that abstracts the user from knowing the complex joins that have to be created to re-assemble the related data into a correct query result.

What if you did not have to deconstruct a database design that inherently has the relationships already defined and matured? What if an end user could simply select a field (attribute) *with the related data from another file* and return the correct results? It is common practice for MV developers to accomplish this by leveraging correlates and translates. Ultimately, a native approach is far more performant than either physically un-nesting the data structures or virtualizing via an SQL mapping approach. It also means *no more flattening of data*.

Translates and links occur at the core level of the database engine thus providing the most efficiency. Any data dictionary definition that is core to the MV database can be utilized, including any data associations. Native access

MV data, as it was meant to be used, is an ideal approach.

This is available with Rocket® CorVu NG for Rocket® U2 and Rocket® D3. CorVu NG provides a robust BI design solution that gives developers the ability to create comprehensive dashboards, which they can use to monitor critical business metrics. The dashboards also allow guided analysis to root cause in addition to visually rich report content using direct native queries to your MV database.

Sure, this all sounds great. You simply point an enabled BI solution directly to your MV database and you are golden. Your users will most certainly be happy. Or will they?

Just because you can now directly access your application's data doesn't

mean you should do so without proper planning. You need a better understanding of the impact caused by users doing BI consumption.

The main reasons for developing data warehouses are: (a) to centralize heterogeneous data from across the enterprise; (b) to offload heavy queries from the transaction systems; and (c) to optimize the data structure and schema for queries versus transactions.

Data is typically pre-aggregated for metric/KPI (Key Performance Indicator), trend analysis for dashboards, and high-level analysis with the ability to then drill down to subsequent lower levels of detail. Unlike the ideal for the transaction system, in warehousing the data columns (fields) are *heavily* indexed for fast retrieval. Since this solution makes the warehouse and the

When was the last time you read something fun?

HDWP BOOKS

Bottom line: Tiago and the Masterless is a book that only a programmer could have written, but not one that only a programmer can love.
- Jon Frater's review on Amazon.com

Upcoming
Tiago in the Tunnels of Krall
Tiago faces the Abyss

<http://www.hdwpbooks.com/books/tiago>

transaction system one-in-the-same, indexes on the system need to be prudently balanced. Transaction write times can be dramatically impacted if too many indexes are maintained. We can't sacrifice the speed of business to improve the speed of analysis.

The Higher Value, Multi-Purpose Approach

The point that I am driving at is that BI activities are different from normal transaction workflows, and can negatively impact the production transaction system's performance. It is well advised to consider off-loading the BI activity to a secondary server. Unlike the typical data warehouse approach, a great way to accomplish this with the U2 and D3 databases is to leverage HADR (U2) or Hot Backup (D3) solutions. The distinct advantage of this

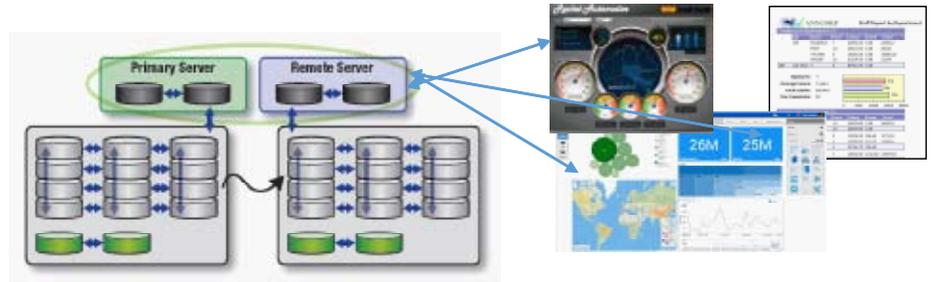


Fig. 1

approach is that by adding these solutions, which are focused on helping prevent disaster, it can also serve a dual purpose as the primary BI data source.

If you are running a U2 database, there is also a Report Server solution that allows you to replicate only the data required for BI analysis; it is not intended for complete duplication or fail-over for the primary system. This works because we don't need to transform the data to another database format.

The data replication is executed dynamically and rapidly. The subscribing server has very low latency compared to the transaction system. This results in the ability for close to real-time analysis without querying the primary system. It does not preclude you from doing some special "real-time" queries against the primary system should the need arise. (Figure 1)

As part of the replication process, you can use write triggers to duplicate selected data to an account specifically for BI use. This allows you to create critical indexes that would not be in conflict with the transaction system if the secondary system needed to switch roles. Additionally, you could also create aggregates in the BI account for higher level metrics and KPIs to optimize query performance for dashboards.

If you want to learn more about Business Intelligence Best Practices for a MultiValue deployment, I will be giving a session focused on Dashboard Best Practices and a Lab on Business Intelligence at MultiValue University.

Rocket MultiValue University 2015 is October 5-8 at Harrah's Hotel and Casino in Las Vegas. To learn more and to register, visit www.rocketsoftware.com/mvu. **IS**

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

FROM THE PRESS ROOM



AccuTerm Mobile v1.0.6 Available

AccuTerm Mobile® allows users to access their MultiValue systems from any mobile device running Apple iOS or Android. The app is optimized for connecting to MultiValue databases but can connect to any host server running Telnet or SSH.

AccuTerm Mobile provides industry leading terminal emulation for ADDS, Wyse & VT terminals, and offers integrated mobile features such as browser, e-mail, and phone functionality for use with legacy applications. Users can take advantage of both SSH and telnet connectivity, allowing you to connect to any MultiValue, Linux, Unix or AIX host system on the go, including jBase, OpenQM, D3, UniVerse, and UniData.

New Features in version 1.0.6 include:

- Added setting to select terminal screen font size, including an option for auto-sized font that scales to fill the screen width with the appropriate number of columns.
- Added support for AccuTerm mouse click reporting which works the same as the desktop version of AccuTerm,

except there is no "hand" cursor (because there is no mouse). Only single tap (click) is supported, and it is reported as the left button. See the AccuTerm 7 help, AccuTerm 7 Reference chapter, AccuTerm Programming section for details.

Bug Fixes:

- Fixed SSH bug which did not recognize any SSH server configured to support both SSH v1 & v2.
- Fixed problem when using Bluetooth keyboard on iPad, where lower half of terminal screen was unused.
- Fixed Android 4.1, 4.2, & 4.3 problem where soft keyboard would automatically close when the Enter key was pressed.
- Fixed problem related to programming function keys using AccuTerm private ESC sequence.
- Fixed problem with underline & blinking effects not always working.
- Minor Wyse & ADDS emulation bug fixes.
- Fixed minor VT visual attribute problem where wrong colors were used if the attribute was reset to normal in a separate ESC sequence from the attribute being set.
- Changed terminal type string sent to host for Viewpoint A2 Enhanced emulation from 'wyse-vp' to 'viewpoint' because many older Linux versions do not recognize 'wyse-vp'.

Changed QM terminal type string for Viewpoint A2 Enhanced from 'vpa2-at' to 'adds4000-at' to match the string used in the desktop version.

- Removed Android file access permission requirement.

If you have purchased AccuTerm Mobile, this update is free. ■



QM Release Notification (3.4-2)

Release 3.4-2 of the QM Multivalue Database is now available for download from www.openqm-zumasys.com/downloads or www.openqm.com/downloads.htm.

This release contains minor enhancements requested by our users:

- The IPADDR configuration parameter can be used to make QM listen for incoming telnet or QMClient sessions on a specific IP address. This may be useful in situations where multiple separate instances of QM run on the same server.
- The QUERY.SORT.NO.CASE mode of the OPTION command makes all SORTs within the query processor case insensitive.

- The ADMIN.USER command now allows selection of whether or not a forced account name should be applied to phantom processes.
- The CREATE.FILE command now has a MULTIFILE option which causes it to create the data portion of the file as a multifile even if no subfile name is given.
- The LIST.INTER, LIST.UNION, LIST.DIFF, and MERGE.LIST commands have been modified to give a significant performance improvement when merging large lists.
- The TANDEM command can be used by administrative users to monitor another QM session. Because this could be a serious security risk, it must be enabled in the QM configuration parameters.
- The query processor COUNT and SELECT commands now support a REPORTING option to display progress information when using a selection clause.
- The PARTIAL.EXPRESSIONS setting of the \$MODE compiler directive modifies the way in which logical expressions are evaluated, such that remaining elements of the expression are not evaluated if they cannot change the result. Without this option, an expression is always fully

FROM THE PRESS ROOM

evaluated.

- The PICK.READNEXT setting of the \$MODE compiler directive causes the READNEXT statement to leave the record id variable unchanged when taking the ELSE clause.
- All variants of the QMBasic LOCATE statement now support an option to force case sensitivity/insensitivity.
- The QMBasic MERGELIST() function can be used to merge two field mark delimited sorted lists.
- Use of SYSTEM(1071) after a successful OSREAD operation returns a dynamic array of file information similar to that returned by the STATUS statement. ■



Pavuk Systems Announces Pavuk Business Platform 4.0

Pavuk Systems today announced the upcoming release of Pavuk Business Platform (BP) version 4.0. Pavuk BP is a comprehensive development and runtime environment for business applications – applications with a modern, Web 2.0 user interface. With many built in business objects, even non-

programmers can create robust database applications.

Pavuk BP 4.0 enables businesses to have best-of-breed applications without needing an army of programmers. Business-specific functions, when needed, are written in a high-level BASIC language. Applications written in Pavuk BP 4.0 run natively in most modern browsers, on both desktop and tablet devices. A mobile interface is scheduled for 2016.

Built on top of the amazing OpenQM database engine, Pavuk BP 4.0 is amazingly fast and scalable from 10 to over 500 concurrent users. Security and audit compliance are built into the system core. Kendo-UI components create a facile and rich user experience that is consistent, ensuring ease of training and adoption by end-users.

“We are extremely pleased with this version. The support of our international partners: Ladybridge Systems, Ltd. and Telerik, have been top-notch,” said Bill Crowell, Pavuk Systems’ CEO and Lead Über Geek. “Versions 1-3 were very good, but version 4.0 is truly amazing. With more great things planned, we look forward to exceeding our customers’ expectations. Our technology is both revolutionary and disruptive. The response of our non-geek beta testers has been overwhelmingly positive.”

Pavuk BP 4.0 supports the

OpenQM database natively. Oracle and MySQL support planned for 1Q 2016.

About Pavuk

Based in Pearland, Texas, USA, Pavuk Systems is a technology firm founded by Bill Crowell in 2007. Mr. Crowell has 36 years of professional experience in developing database applications in various industries including health care, accounting and energy. Pavuk Systems operates on a distributed development model with an alliance of programmers in the USA, UK, Slovakia, and Bulgaria.

Contact: Pavuk Systems, LLC, 2716 Shauntel St, Pearland, TX 77581 USA, or phone: +1 713.589.9711 or e-mail: bcrowell@pavuk.com ■



Rocket Software Introduces the Rocket Discover Business Intelligence Solution

Rocket Software is pleased to announce the launch of Rocket Discover, a self-service, intuitive data preparation and discovery solution to give business managers and executives the speed and confidence to make key

decisions. Rocket Discover allows users to easily access, manipulate, prepare and visualize data to expose insights and help them make critical business decisions that will drive growth and outpace the competition.

“Companies facing critical business issues know that they need to act fast and act smart, but they often don’t have the right information at their fingertips to investigate and quickly resolve the problem,” says Bob Potter, Senior Vice President and General Manager of Business Information and Analytics at Rocket Software. “While there are many self-service BI solutions on the market, we continue to hear about critical functionality that these solutions lack: self-service data preparation, access to the mainframe, and scalable performance. With this in mind, we built Rocket Discover to be a world-class visual data discovery tool for business users that addressed these gaps.”

Now business users can easily access and enrich data from multiple sources – ranging from mainframes to databases to the cloud. With self-service data preparation and discovery in one solution, they can get business insight within minutes, not days or weeks. For optimal performance, Rocket Discover was engineered with a data-retrieval architecture that combines in-memory

FROM THE PRESS ROOM

performance with live query access; a hybrid approach that can handle today's high volume thresholds.

Today's business professionals know their data, their business challenges, and the need to quickly connect the two to address these challenges. Rocket Discover's emphasis on intuitive design for self-service data preparation and discovery makes this possible. Users can easily interact within a rendering canvas to create visual data process flows and dashboards through simple drag-and-drop movements. And when it's time to collaborate, team members simply drag a dashboard into the chat window for immediate dialogue on strategy and opportunities. Rocket Software has specialized in database connectivity, preparation and solving customer problems for over 25 years, which is why Rocket Discover is differentiated with:

- Connectivity to disparate data that includes IBM z Systems and IBM Cognos TM1
- Self-service data preparation and data discovery in one solution
- Data Lineage for governance
- High Performance data access for rapid decision making

Jan Baesjou of ABM Systems, a Rocket Strategic Partner, says that Rocket Discover is redefining BI. "We have used

Rocket's BI products for over 20 years, and Rocket Discover is the perfect addition to our business planning arsenal. There are plenty of BI tools out there, but most of them focus mainly on dashboarding. Rocket Discover has capitalized on their data preparation IP in which they've excelled for many years. Having simplified data preparation within a BI solution is a key feature and it is a win for our business users."

"We have been using Rocket Discover for several months, and are impressed with the potential that the application presents," says Rachel Byrom from Oxford City Council. "I come from the IT side of the business, so I'm excited to get Rocket Discover into the hands of as many of our business users as possible. Rocket Discover will allow staff to directly access data and create dashboards that are contextual to their needs, when and how they need it. This will ultimately reduce the number of requests to my team for reports and dashboards, allowing us to focus more on development projects and system improvements."

Subsequent releases of Rocket Discover will continue to capitalize on Rocket Software core competencies by providing additional features in the areas of data connectivity, maximizing the value of the mainframe, increased levels of intelligent guidance during data preparation, and enhanced collaboration.

About Rocket Software

Rocket Software is a global software development firm that builds enterprise products and delivers enterprise solutions in the following segments: Business Information and Analytics; Storage, Networks, and Compliance; Application Development, Integration, and Modernization; and Database Servers and Tools. Rocket is engaged in business and technology partnerships with IBM, EMC, Fujitsu, Hitachi Data Systems, Avaya, Epicor, Infor, and many others. The company is headquartered in Waltham, Massachusetts, USA. ■



Rocket Software is Proud to Announce Rocket U2 Web DE 5.2.0

Rocket® U2 Web Development Environment (U2 Web DE) provides a standard way to create objects (Rocket U2 data and business logic) for use in your web and mobile applications.

Monitoring enhancements include features such as a traffic light indicator to illustrate the system's performance statistics, response

time tracking, license consumption, and better exception reporting. Along with the logging improvements, this information helps you to stay on top of the health of your system.

And now you can expose your existing web logic for use in mobile applications via RESTful services. We are ready to help you securely take your operations to the next level with Rocket U2 Web DE 5.2.0!

Highlights of this release include:

- Monitoring – Keep an eye on your website responsiveness via traffic light indicators, troubleshoot when faced with high demand, and identify areas to improve or adjust license consumption.
- Logging Improvements – Quickly identify slow processes, keep track of the distribution of requests across database licenses, and see any database or basic errors that may have occurred. Logs can be set to gather statistics on an adjustable, timed basis and can be turned on or off without restarting the scheduler.
- RESTful Services – Expose your existing web logic and reuse your existing Rocket U2 Web DE RedBack Objects (RBOs) in other applications and processes via RESTful services. ■

ROLL YOUR OWN (HIGHLY EXPERIMENTAL) MULTIVALUE DATABASE PART 2

In this series of articles, we are going to build a brand new flavor of MultiValue database engine. I call it Lumpy, because data isn't flat. Also, because programmers shouldn't be allowed to name things. While Part 1 laid the groundwork, a database isn't a database if you can't CRUD (Create, Read, Update, Delete) it up.

Remember, while Lumpy is implemented in PHP, it could be implemented in any language. The goal is to build a roadmap of features and create a reference set of working commands in one language. Where you take it is up to you.

"Oh, Pen," Says Me

Since Lumpy is MultiValue, we need to start any discussion of file transactions with the OPEN statement. OPEN establishes a handle — which can be used by READ, WRITE, and DELETE — to specify which file is responding to the actions.

Since the language isn't BASIC, we have some hoops to jump through. Our variables won't be their variables, but this is only useful if you can easily exchange values between the host language and the Lumpy commands. To accomplish this, before we can dig

The goal is to build a roadmap of features and create a reference set of working commands in one language. Where you take it is up to you.

into the CRUD, we need a way move things back and forth. Here's an example of our get and set commands:

```
Lumpy->command('set
attendee.id',phpVariable);

phpVariable = Lumpy-
>command('get attendee.
rec');
```

We will treat the file variable, the record id, and the record itself, as Lumpy variables. OPEN will look like this:

```
Lumpy->command('open
attendee to attendee.
handle');
```

The OPEN statement should allow for THEN and/or ELSE actions:

```
If (Lumpy->command('open
attendee to attendee.
handle')['error'] = false)
{
}
else
{
}
```

All of this will be extended into PHP using the verbs.php module we started developing in the last issue. Lumpy programs will start with this PHP command:

```
include_once('./verbs.
php');
```

The code for verbs.php can be found at hdwp.com/lumpy/

Reed and Rite

Now that we've laid the groundwork, we can get to the core commands for this article. Here's an example of the syntax for the Lumpy version of READ with a THEN clause:

```
Lumpy->command('set
attendee.id',phpVariable);

If (Lumpy->command('read
attendee.rec from
attendee.handle, attendee.
id')['error'] = false)
{

phpVariable = Lumpy-
>command('get attendee.
rec');
}
```

The code for verbs.php can be found at hdwp.com/lumpy/

Of course, we don't have just one READ command in MultiValue. *Spoiler: not all of them will make an appearance in our new system.* Lumpy will

support READ and READV. It won't support READT (tape), READU (locking), and READVU (also locking). Tape processing and locking are both out-of-scope. Since all the code will be available, you are welcome to provide those extensions.

Lumpy will also skip READSEQ (foreign files) and MATREAD (dimensioned arrays). Don't feel the need to add these two on your own time. MATREAD and all its related commands are omitted because fixed arrays don't have the same performance differential in PHP that they would have in other MultiValue implementations. MAT is useless in Lumpy. SEQ-style commands... we'll get to that next issue. While we *are* leaving them out, we *aren't* really leaving them out. It's complicated because it will be easier. You'll see.

Part of the fun of roll-your-own is customization. While that will mean cutting some commands, as we've already discussed, we can also jazz up Lumpy by adding new commands. READX tests if the record exists. It will not return the content, just a true/false. You could do this with a READ or READV, but READX will be more efficient.

Some examples of our three READs:

```
include_once('./verbs.
```

```
php');
```

```
Lumpy->command('read
attendee.rec from
attendee.handle, attendee.
id');
```

```
Lumpy->command('readv
attendee.rec from
attendee.handle, attendee.
id, -1');
```

```
Lumpy->command('readx from
attendee.handle, attendee.
id');
```

The code for verbs.php can be found at hdwp.com/lumpy/

Reading only works if there's content. The WRITE command is our mechanism for adding data to our database. As with READ, WRITE will lose some variants. It will also gain two new ones. Say goodbye to WRITET (tape), WRITEU (locking), and MATWRITE (dimensioned arrays). Say hello to WRITEC and WRITER. WRITEC is a specialized version of WRITE which will create new records but not replace existing records. WRITER is the reverse. It will (update) replace existing records only, but never create new ones.

DELETE comes through as-is.

```
Lumpy->command('write
attendee.rec on attendee.
handle, attendee.id');
```

```
Lumpy->command('writev
attendee.rec on attendee.
```

```
handle, attendee.id, -1');
```

```
Lumpy->command('writec
attendee.rec on attendee.
handle, attendee.id');
```

```
Lumpy->command('writer
attendee.rec on attendee.
handle, attendee.id');
```

```
Lumpy->command('delete
attendee.handle, attendee.
id');
```

The code for verbs.php can be found at hdwp.com/lumpy/

Delete, Not to be Confused with Delete

MultiValue has a second form of DELETE, but it does not deal with file I/O so it was omitted from this article. READNEXT does deal with file I/O but not in the same way as other READ commands. READNEXT has a featured role in Part 3.

Verde Ick

Part 1 established the commands for creating databases, accounts, and tables. This part was all about CRUD, Let's see how we did.

CRUD scorecard for Lumpy:

- Create: WRITEC
- Read: READ, READV
- Update: WRITER
- Delete: DELETE

Plus, we still have the traditional WRITE, WRITEV, and the new READX.

Hide and SEQ

While we've met our goals, the Lumpy odyssey is not over. Next up, we will discuss foreign data structures, dictionaries, and LQL (Lumpy Query Language). **IS**



Informer™
Quick and Easy Business Intelligence

The BI Survey 14 ranks Entrinsik Informer #1 in "Ease of Use" and "Chosen as Standard" categories among dashboarding & ad-hoc reporting vendors.

Find out why at www.entrinsik.com

MARKETPLACE

ACCOUNTING

Natec Systems

www.natecsystems.com | nrector@natecsystems.com

mvQB 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

Phone: 303.465.9616
E-mail: mvqb@natecsystems.com
Website: www.natecsystems.com

NATEC Systems
Providing Solutions to your MultiValue Questions

COMPLIANCE

SJ+ Systems Associates

www.sjplus.com | sjoslyn@sjplus.com

CONSULTING

Drexel Management Service

www.drexelmgmt.com | dconboy@drexelmgmt.com

Execu-Sys, LTD

www.eslny.com | mh@eslny.com

HDWP

www.HDWP.com | results@HDWP.com

J Squared Software, LLC

www.jsquaredsoftware.com

30+ years experience Business Applications & Support
MultiValue platforms - per diem or fixed price
Juddlove@jsquaredsoftware.com 203-325-0275

Modern MultiValue, LLC

www.ModernMultiValue.com | info@ModernMultiValue.com

PICK Programmers Shop

www.pickprogram.com | brian@pickprogram.com

Precision Solutions

www.precisionline.com | Kevin@PrecisOnline.com

DATABASE

Ladybridge Systems Ltd

www.ladybridge.com | sales@Ladybridge.com

REPORTING

Brian Leach Consulting, LTD

www.brianleach.co.uk | brian@brianleach.co.uk

Entrinsik

www.entrinsik.com | sales@entrinsik.com

TERMINAL EMULATOR

Zumasys

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

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

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

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

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

MultiValue Web Connector

Part 1

BY KEVIN KING

If you've been following these articles for the past few months, you may be sensing that I'm a bit stoked about all that's happening in the world of web technology. And why not? With new versions of HTML, CSS, and Javascript fresh out of the oven, plus an abundance of new web-enabled tablets and phones hitting the market, there is truly MUCH to be excited about!

All the cool web technology available right now doesn't mean much if we can't make it work with our MultiValue systems, however. Our information, the valuable lifeblood of our companies, lives in those systems, so if the web can't talk to our databases then ... why bother, right? Well my friends, prepare yourself to bother because I'm about to show you that a browser CAN talk to MultiValue databases, and it's not nearly as difficult or expensive as you might have been led to believe.

Before we get too far down this rabbit hole, allow me to first acknowledge that there are a number of vendors who have spent a good deal of time, money, and effort creating solutions

A browser CAN talk to MultiValue databases, and it's not nearly as difficult or expensive as you might have been led to believe.

to make the web accessible to MultiValue programmers. My point here is not to discount those efforts or to undermine them in any way. Rather, for those who don't mind a bit of bare metal programming and diving deep into the dirty bits, I'd like offer you something interesting to think about, something fun to "research," and if all goes to plan, this exercise might just shift your perception a little.

In an effort to keep this simple, we'll be building this connector as a series of incremental steps. Today we'll create the MultiValue parts needed for our connector. Later we'll create some code to connect our MultiValue solution to a web server. Finally, we'll tie it all up with one more script that will give us secure and controlled access

to our MultiValue data and programming from a typical website.

Ready to jump into the dirty bits? Let's do it!

Most MultiValue systems have a shell program that is launched from the host operating system command line. For example, Unidata has "udt", QM has "qm", Universe has "uv", etc. While these shell programs are usually launched from a login script when someone makes a telnet connection, these commands can just as easily be launched from an ssh session, cron job, or even from a web server. This is this foundation on which we will build our connector.

On the MultiValue side, our connector is nothing more than a small Basic program that accepts some input and does something practical. Of course, "something practical" could mean just about anything, so we'll build our routine as a simple wrapper that can call a subroutine. With this foundation in place, your team of geniuses can create subroutines for whatever functionality your web solution may require.

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>



Before we create this program, however, we need to figure out how to ask our connector to do something. Most MultiValue systems provide some way of reading environment variables from the underlying operating system, so let's start there. We'll set an environment variable called WEBREQ with a request, start up our MultiValue database, and then grab the response when the shell finishes.

Another thing to consider is how we will launch our Basic routine. Most systems have some way of starting the shell and launching a specific program, but they don't all work the same way. Therefore for portability we'll leverage the login procedure (named "LOGIN" on Unidata or may be named accordingly to the account on other platforms) to check for our environment variable and do something accordingly. If our environment variable is not set (as would be the case for a normal telnet connection) our program can stop and allow the login procedure to continue normally.

Figure 1 shows how this might be accomplished using a program named RUN.WEB.REQUEST. This program was written for Unidata but could be easily translated to a number of MultiValue platforms.

In this program we read the environment variable WEBREQ and expect it to contain three comma-delimited pieces: 1) the subroutine to be called, 2) the name of a file where the output will be written, and 3) the input data for the routine. Once we've parsed the three parameters from the environment variable, we use a CALL @ to call the subroutine as named and then write whatever result was returned to a directory file named WEBOUT (which I usually map to the _HOLD_ directory on Unidata or \$HOLD on QM).

Despite the simplicity of this routine, there are many details to explore. For example, you might wonder why we're passing the input in an environment variable but writing the output to disk. Good question!

The input from the web might contain confidential information (i.e. credit card numbers) so we don't want that information EVER written to disk. Of course, with the output written to disk — the only portable way to get information out of the database and into the underlying operating system where it can be picked up by an external program — we must be careful to avoid including sensitive information in our output or find another way to get the information out of the database shell and into the calling operating system.

(Printing the output to the screen and capturing it in a calling script can work too. However, I like to use CRTs and COMO for debugging and having debugging output mixed with the real output data is a little too inflexible for my tastes. Also note that this routine

```
*
* Program Name: RUN.WEB.REQUEST
* Written By:   Kevin King
* Project:     Spectrum Demo
* Date:       26 Nov 2011
* Description: This routine will look for an
* environment * variable WEBREQ and will call a
* subroutine as requested when found.
*
WEBREQ = GETENV('WEBREQ')
IF (WEBREQ NE '') THEN
  OPEN 'WEBOUT' TO F.WEBOUT THEN
    SUBR.NAME = FIELD(WEBREQ, ',', 1)
    OUT.NAME  = FIELD(WEBREQ, ',', 2)
    DATA.IN  = WEBREQ[COL2() + 1, 999999]
    DATA.OUT = ''
    CALL @SUBR.NAME (DATA.IN, DATA.OUT)
    WRITE DATA.OUT ON F.WEBOUT, OUT.NAME
    CHAIN 'OFF'
  END
END
*
STOP
END
```

Fig. 1

leaves basic error handling as an exercise for the reader.)

Rather than only calling subroutines, we could expand upon this concept to include fundamental database operations like reading, writing, deleting, returning a select list, executing a TCL command, etc. However, because these operations necessitate more advanced security considerations, we will be focusing only on calling subroutines. These subroutines can then be written as securely or insecurely as you prefer.

Speaking of security, you've probably also noticed that our routine allows a subroutine to be called with absolutely zero authentication or additional security. Many MultiValue systems don't ask for additional authentication once someone has logged in to the box, which may be fine for telnet users, but

```
SUBROUTINE TEST.SUB (DATA.IN, DATA.OUT)
*
DATA.OUT = OCONV (DATA.IN, 'MCL')
*
RETURN
```

Fig 2

should be given additional consideration for web-based connections.

Figure 2 is a very brief example of a subroutine that might be called from our connector. Note that this subroutine takes two formal arguments, one for input and the other for output. In this simple example we're converting the input string to lower case, but as was mentioned earlier, this subroutine could do absolutely anything it wants with the input in order to produce the output.

With nothing more than these two routines and a minuscule edit to the login procedure to call our starting

program, we now have a way of invoking a MultiValue subroutine from the outside world. Next time we'll take a look at how we can use the open source Apache web server and a remarkably small bit of PHP to provide this functionality to a larger audience! **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.

With nothing more than these two routines and a minuscule edit to the login procedure to call our starting

eXtra Newsletter

Stay on top of Industry News

- ◆ Tech Tips
- ◆ Job Postings
- ◆ New Products
- ◆ Corporate Updates



www.intl-spectrum.com/newsletter