

\$7.00 U.S.

INSIDE: MultiValue in the Cloud

INTERNATIONAL Spectrum®

THE MULTIVALUE  TECHNOLOGY MAGAZINE | SEP/OCT 2011

PLUS!
Both Text and
HTML: Multi-
part E-mail

Are

Persistent Connections

a Thing of the

Past?

intl-spectrum.com



Advanced database technology for **breakthrough applications**



This makes applications fly.

Embed our post-relational database if you want your next application to have breakthrough features, run with blazing speed, be massively scalable and require minimal administration.

InterSystems **Caché**® has advanced object technology that makes it easier to build applications with XML, Web services, AJAX, Java, and .NET. And Caché can run SQL up to 5 times faster than relational databases.

With its unique Unified Data Architecture™,

Caché eliminates the need for object-relational mapping. Which can reduce your development cycle by as much as 40%.

Caché is available for all major platforms – and it supports MultiValue development. Caché is deployed on more than 100,000 systems worldwide, ranging from two to over 50,000 users.

For over 30 years, we've provided advanced technologies for breakthrough applications.

INTERSYSTEMS

Download a free, fully functional, no-time-limit copy of Caché, or request it on DVD, at InterSystems.com/Advanced6WW



FEATURES | SEPTEMBER/OCTOBER 2011

6 **MultiValue Communications: The Persnickety Persistence**

Problem For decades, MultiValue applications have followed the model of having an always active, persistent connection to the server. This worked well when we had control over the connection, primarily by using a dedicated cable to a serial RS-232 port physically located on a board in the machine. Then came networks, where everything is dynamic and out of our control. Yet in 2011, how many of our MultiValue application still stick to trying to emulate the old persistent connection model of the last century? **BY KEVIN KING**

8 **Useful Tips on Migrating from Legacy PICK** Your system sends up certain red flags when its performance has degraded. These signals might include missing your maintenance windows, users complaining about slow system response times, or unreasonably high system maintenance fees. Perhaps your customers are saying the application needs modernization, but the programming backlog is six months and growing. How do you solve these problems? Replace your legacy PICK system! **BY DR. JOHN F. NUNZIATO, U2 DATA AND APPLICATION MIGRATION SPECIALIST, ROCKET SOFTWARE**

12 **Sending E-mail from MultiValue Programs - Part 5: Creating and Sending E-mails with HTML and Text Sections** In the early days of e-mail, text was King. If you dared use that horrible, evil, Microsoft Outlook mailer and sent HTML-based e-mail, you had better be wearing asbestos underwear, because you were going to get flamed. Today, HTML e-mail is the norm for all large businesses (at least in the USA), and the text-only pundits are considered to be somewhat Luddite. Then came mobile devices with small, hard to read screen sizes. What to do? What to do? Enter the realm of multipart html/text e-mail. **BY NATHAN RECTOR**

22 **Business Tech: MultiValue in the Clouds** Although there are still concerns with Cloud Computing — data privacy, guaranteed service levels, etc. — it appears that the Cloud is here to stay. So what will it take to be able to deploy your MultiValue applications in the Cloud? (Hint — you are almost there already.) **BY CHARLES BAROUCH**

24 **Auditing Database Changes with UniVerse Indexing Subroutines** UniVerse has file triggers which can be used to call a subroutine whenever a new record is added, changed, or deleted from a file so that the change to the database can be recorded in an audit trail for IT governance and compliance requirements. But full blown triggers come with a certain amount of performance overhead. For simple auditing, using indexing subroutines with a new system variable available at release 11.1 is a lighter weight alternative. **BY CLIFTON OLIVER**

DEPARTMENTS

From the Inside page 4

From the Press Room page 16

Tech Tip page 28

Clif Notes page 31

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

BY NATHAN RECTOR



I need your help in getting the message out that MultiValue is a modern, robust, solution-oriented platform that works well in a multi-system, mainstream environment.

Are you interested?

I've asked many people to write articles over the past few years, and I have been rewarded with some very interesting articles. Most of these authors are techies like you, not professional writers. The best articles are the ones provided by the end users in the marketplace.

I understand that writing can seem daunting, and you wonder where to find the time to do it. Well, we are here to help.

The most common problems that new authors have are:

- Don't know what to write about
- Company restrictions on propriety code
- Not a writer and don't know how to create a finished article

Issue: Don't know what to write about

I can understand this one. I have to come up with a topic for the "From the Inside" every issue. It seems like this is the biggest issue, but in reality, it is the easiest to address. The answer: Call me or e-mail me and ask.

I have a whole list of topics that I would like to see articles written about, but most of time I want to hear about the solutions you have developed for your business.

If you are reading this and would like to see a specific article in International Spectrum magazine, please feel free to contact me and let me know. I can add it to my growing list of articles, webinars, and conference topics.

Issue: I'm not a writer

Well, truth be told, neither am I. My spelling is horrible. My grammar makes Clif pull his hair out. (Editor's Note: I plan to bill Spectrum for a hairpiece.) And about every article I write, he has to clarify at least one paragraph because he's not quite sure what I'm trying to say.

My suggestion is, don't try to be a writer. Be yourself. Write the article as if you are talking to someone or making notes on how to do something. We can take that rough material and make it a finished product.

The article doesn't have to be a Word document. It can be typed out in e-mail. Take an hour for lunch and use your smart phone to write the article.

So, if you are worried about your writing — please don't be. Let us address the polish of your articles. Trust me, the International Spectrum Staff is very good at doing this.

Issue: Company Restrictions on Proprietary Code

This can be a major hang up for a lot of people, but you don't have to provide your company's code in the article. You'll have to describe the problem you were solving and an outline of the solution.

Many times the article is about a concept or suggested way of doing something, and you can strip the coding examples down to a few lines of code to show examples, without giving away anything your company owns.

One of authors here at International Spectrum has this same problem. He worked with his management, and worked how to write the article to cover the topic without using any of the company code. Most of the code

Continues on page 29

INTERNATIONAL Spectrum

SEPTEMBER/OCTOBER 2011

NATHAN RECTOR
President

CLIFTON OLIVER
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., 8956 Fox Drive #102, Thornton, CO 80260; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com. Copyright 2011 International Spectrum, Inc. All rights reserved. Reproduction in whole or in part, without written permission, is prohibited.

PRINTED IN USA

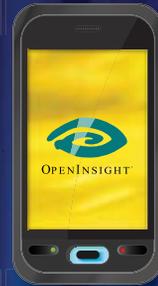
NEWS RELEASES/UNSOLICITED ARTICLES

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.

twitter.com/intlspectrum

intl-spectrum.com/facebook



MOBILE COMPUTING

An essential part of business today is mobility. We store and access crucial data off-site more than ever before.

DATA ENCRYPTION



Your information needs to be accessible and yet still secure.

ARE YOU READY? WE ARE.

We are Revelation Software, creators of OpenInsight Development Suite 9.3, bringing you one of the best mobile computing and data encryption toolkits on the market. Go to revelation.com and take a test drive today.

Revelation Offices – Global

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

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

Revelation Software Australia Pty Ltd.
PO Box 300
Brookvale, NSW 2100
Australia
Phone: +61 2 8003 4199
Fax: +61 2 9332 6099
Email: info@revelationsoftware.com.au
Web: www.revelationsoftware.com.au

REVELATION
SOFTWARE

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

MultiValue Communications

The Persnickety Persistence Problem

Last issue we took a look at the good, the bad, and the ugly of the Telnet protocol in the context of our MultiValue applications. In brief, the Telnet spec is nearly 30 years old, the terminal types emulated over Telnet were antiques years ago, the protocol is implicitly insecure, and if that's not enough, it drops connections far too frequently.

We also briefly mentioned SSH and Telnet over SSL, and while these can provide improvements in security, they do little to nothing to improve the reliability of the connection. With miles of cable and countless devices in the network that are out of our control, improving reliability of persistent connections is problematic. But consider this: Might we **improve** reliability simply by **decreasing** our dependence upon persistent network connections?

On behalf of an industry that has been dependent upon Telnet from the be-

So if we can't stop these failures, perhaps we should instead be focused on minimizing their impact.

ginning, the heresy of such a radical recommendation is well appreciated. However, as we look another 30 years into the future, perhaps radical thinking is exactly what we need, and now is exactly when we need it.

If we're perfectly honest with ourselves, we have to admit that all of our best planning will not prevent someone with a backhoe inadvertently digging a hole through the wires that are carrying all of our precious communications. It happens. Moreover, networking equipment fails and with all these devices between us and our information, something is definitely going to fail at some time, and it'll likely happen at the least convenient moment.

So if we can't stop these kinds of failures, perhaps we should instead be focused on minimizing their impact. To achieve that, let's start by taking a brief

look back at our history of persistent connections.

Though it may seem obvious, we remain tethered to persistent connections simply because of how we communicate with our MultiValue applications using the Telnet protocol. Telnet is a persistent protocol; as long as there's a channel of communication between the client and server, Telnet works very well. Break that channel, however, and things can get ugly. Connections are dropped, sessions are terminated, and whatever was in progress at the time ceases to be in progress (barring some form of pre-established intervention).

Equally interesting, Telnet is a single channel protocol. In other words, when you establish a Telnet connection to your MultiValue system, you establish each connection one at a time. If that connection drops for any reason, you can certainly start another connection but there is no provision in the Telnet protocol for redundant or re-connectable connections to protect against interruptions and disconnects. Therefore, and at the risk of sounding alarmist, with Telnet we are all just one



Digital Version
On-Line

Scan Me!

blip away from getting disconnected at the worst possible moment. For some, this may rarely happen. For others of us who seem to always be hundreds of miles away from our systems, these kinds of disconnects happen far too frequently.

Persistence has become so common throughout the MultiValue world, I would venture to say that most don't think much about it until the network hiccups and a connection is lost. THEN we think about it as we establish a new connection, clear any rogue record locks, check to be sure nothing was damaged or left hanging by the disconnect, and ... possibly mutter a few questionable words throughout the process!

With this in mind, do we REALLY need persistence? Is it an absolute necessity that we keep a live connection to our MultiValue servers all the time in order to maintain a dialogue with our applications? Or do we simply continue to rely on persistent communications because this is just a fact of Telnet life, and that's the way things have always been done?

The greatest advantage of persistence is, of course, state management. With a persistent connection to a back-end server there is a block of memory available where we can store all sorts of information. Records can be read from disk into memory, manipulated in memory for the entire length of the connection, and written out as the application sees fit. As long as nothing bad happens to the network, it's all a beautiful thing.

There is of course a cost to all of these connections. Each persistent connection requires some long term commitment of memory, and depending on the number of simultaneous connections and data needed by the application, that can work out to be a pretty big chunk. There's also the ongoing network traffic for maintaining all of these connections, and while Telnet itself is really quite lightweight, anything

multiplied by hundreds or thousands of clients can introduce a concern.

But back to the question: Do we really need persistence? If the greatest advantage that persistent connections offer is state management, and especially where that state is always just one network blip away from oblivion, then there must be a better way. Fortunately there is a better way, and the PHP language offers a stellar example of one possible solution.

In PHP, state management is handled much like MultiValue common variables. Simply set a "session variable" now and when you come back later that variable will still have its assigned value. In a web context where each request to the back-end server requires a separate and independent request, this actually looks to the web application like a block of persistent memory. But alas, it is a cleverly crafted illusion.

Continues on page 20

mits
Advanced Reporting
and Business Intelligence

Announcing...
MITS Discover 7.2

- Full Screen Dashboards
- Prompted Scorecards
- Browser-Based Mobile Device Access

MITS Discover now supports the creation of full-screen dashboards and scorecards. The Executive Dashboard scrutinizes high-level company metrics. A prompted scorecard shows summary metrics for a particular vendor, customer or branch (as examples). Deliver reports and dashboards automatically on a schedule. All summary metrics link directly to an underlying interactive report. View your MITS data from an iPad or Android mobile device.

See MITS Discover 7.2 in action!
Call us today at 888 700 6487
or visit us online at www.mits.com

©2011 Management Information Tools, Inc. All rights reserved.

Useful Tips on Migrating from Legacy PICK

Recognize When to Commit to Migrating

Your system sends up certain red flags when its performance has degraded. These signals might include missing your maintenance windows, users complaining about slow system response times, or unreasonably high system maintenance fees. Perhaps your customers are saying the application needs modernization, but the programming backlog is six (6) months and growing. How do you solve these problems? Replace your legacy PICK system! This article will share the best practices that Rocket U2 has learned through years of experience.

You might be hesitant to undertake a migration. Are you worried about the cost or afraid to commit the time? How can you know how long a migration might take? You might be comparing the amount of suffering your current system is causing versus the pain you might experience while migrating to a new system. You are not alone. Many in the legacy Pick community are facing these pains and questions. Read on to learn about the steps in the migration process, and to get some tips that will help you along the way.

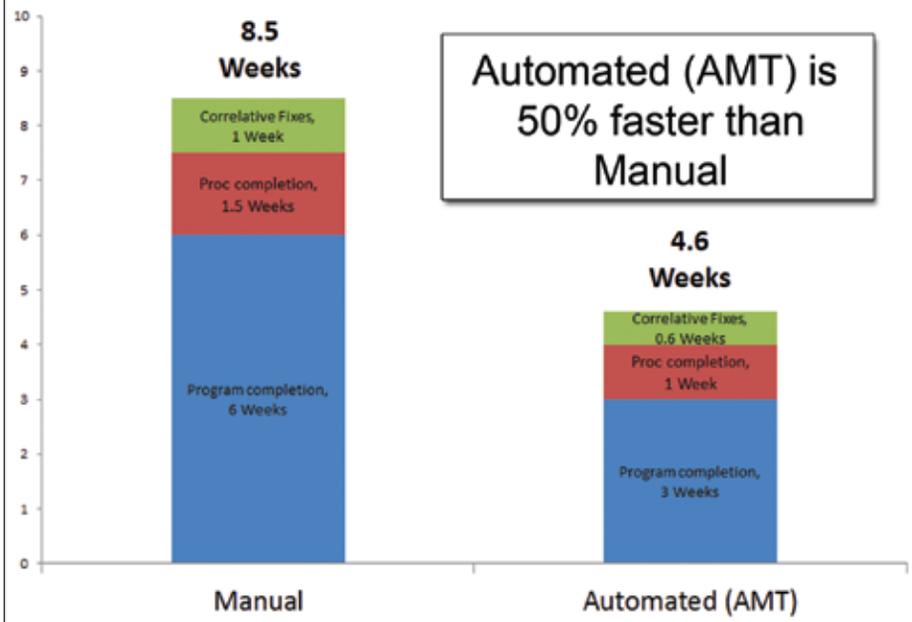
A typical migration involves a thirteen (13) to sixteen (16) week effort for applications with 2,000 programs, 500 PROCs and about 1,500 data files and dictionaries.

Take Inventory of Your Current System and Write Down a Wish List for the Future System

A successful migration starts with good planning. Collect information about

your IT department in order to assess resources and set realistic timelines. Gather a description of IT department personnel; record their experience level and area of expertise. Inventory your current hardware platform; list all system configurations, printers, and communications devices so that you will be aware of the hardware you can reuse or need to purchase for your new system. Remember to document all third party software being used with your system and research if that third party still supports and markets the software. If that third party software is no longer supported, now would be an

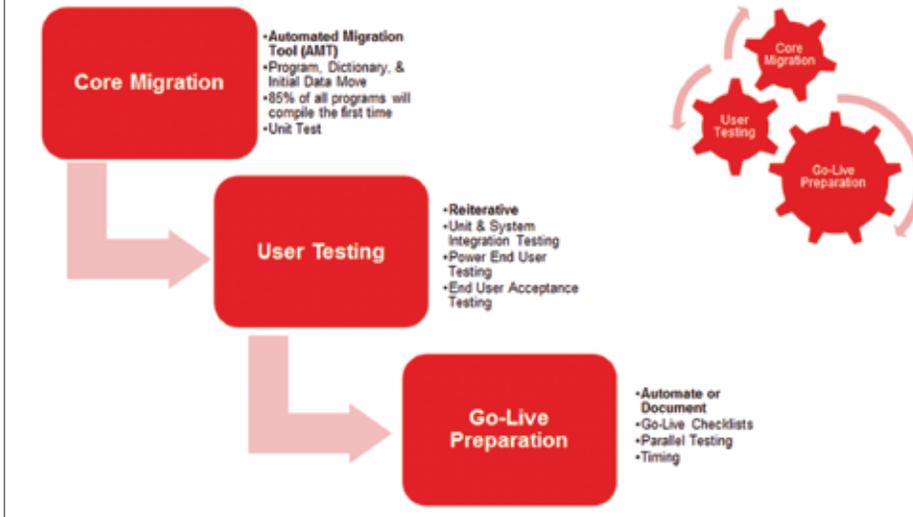
CORE MIGRATION Manual Vs. Automated (AMT)



Digital Version
On-Line

Scan Me!

Application Migration – The Steps



excellent time to consider an improved replacement or “workaround” for that software!

After recording your existing setup, envision your ultimate wish list for software and hardware enhancements for the new platform. Don’t be shy, aim high. You are no longer going to be restricted by an antiquated system! This transition gives you an opportunity to address long needed requirements, such as compliance, High Availability, and software or network modernization.

Next identify the number of program files, PROCS, PARAGRAPHS, data files and dictionary files that are currently being utilized. Although you might find this the most tedious stage described so far – shouldn’t you really be aware of this level of detail about your current system?

Keep It Clean and Neat and Practice Makes Perfect

Clean house by tidying up your application. Sign up your users to be testers during the migration process, and manage expectations during the “Go-Live”. Be sure to emphasize to the people testing the system that they will be held accountable for ensuring every-

thing works properly in the new environment.

Make clear and concise documentation. Craft a checklist and assign each task to the proper resource. Carefully

describe each step so that your staff can reference this documentation at each stage of the process. Rehearse your “Go-Live” as many times as possible, to uncover any issues that might arise before implementation and develop a workaround or resolution to each prior to the “Go-Live” date.

How Long Does the Migration Process Take? How Can I Make It Go More Smoothly?

A typical migration involves a thirteen (13) to sixteen (16) week effort for applications with 2,000 programs, 500 PROCS and about 1,500 data files and dictionaries. The migration is done in stages: the conversion of programs, “Unit Testing” by the technical staff, user testing, pre and post “Go-Live” testing/support, and fixes/workarounds.

Continues on page 10

The Post-Relational Database Reference

Project: UniVerse™

www.keyally.com/prdb

Free resource

Hip
•
Deep
•
In Code



WWW.KEYALLY.COM
RESULTS@KEYALLY.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

1 Penn Plaza, Suite 3306
New York, NY 10119

(212) 967-0505

(800) 423-1964 x 302

Fax: (212)947-8593

Email: mh@eslny.com

Consultants Needed Nationwide

USEFUL TIPS ON MIGRATING FROM LEGACY PICK

Continued from page 9

Choosing the right target platform is your most important decision. You want to preserve your company's investment in business logic, while having minimal impact on your users and production environment. The U2 databases offer the most in terms of compatibility, price, performance, selection of tools for application modernization and more.

With the U2 Migration Factory, Rocket Software facilitates the migration process, allowing it to be as painless and seamless as possible. Rocket U2 has the leading truly MultiValue databases in the industry and has completed over 1,000 migrations worldwide. A free trial migration can be performed as a Proof of Concept (POC) to guide the customer in the planning and conversion stages of the migration process. Rocket U2's Automated Migration Tool

results in a 50% faster Core Migration Process than performing a manual process. This time savings has been proven by performing similarly-sized migrations with and without using the tool.

If you are on old hardware, the clock is ticking. If your legacy system has the business logic you need but is lacking in performance or appearance, a migration can open the door to using modern technology to solve these issues. Following the steps outlined above, you can achieve a successful migration to a more modern database and open the door to a brighter, longer future. **IS**



DR. JOHN NUNZIATO

has many years of experience (over 1000 implementations) optimizing

U2 performance on a wide range of UNIX platforms during migrations and system health checks.

Be prepared for Windows 7. Take the right path with the all-new **AccuTerm 7**!

If we had used **AccuTerm** we'd be there by now.



www.asent.com

U2 Migration Factory

Migrate to the Best MultiValue Databases Available:
UniData® or UniVerse® – Powered by Rocket U2

Business intelligence through U2 DataVu

Integration with enterprise systems
through Web Services

Web enablement using U2 Web DE

Development using .NET or Java

High availability to ensure 24x7
operation of critical business systems

U2 Migration Factory

For more information, email: U2Services@rs.com
or visit: <http://www.rocketsoftware.com/u2>

Rocket

U2®

Powering Business Solutions

Rocket U2 comprises UniData® and UniVerse®, two extended relational databases with integrated development environments designed for rapid, cost-effective vertical application development.

The products are embedded in solutions provided by our partners and used by thousands of organizations from small to medium businesses (SMB) to large enterprises.



Sending E-mail from Your MultiValue Programs

Part 5: Creating and Sending E-mails with HTML and Text Sections

Now that I have covered the basics of sending an e-mail from your MultiValue system, let's look at making the e-mail look pretty. A common question is, "how can I present my e-mailed invoices as HTML, or send a newsletter, or add a nice signature to my e-mail?"

Personally, most e-mails I send are text-only, unless I'm trying to get someone's attention, or it is better to present something in a prettier format. If you have ever looked at the International Spectrum Newsletter, there are two sections to it: HTML and Text. That way if you have preference of one over the other, or you are displaying the e-mail on a mobile device, it still displays nicely.

To add both HTML and Text to your e-mail body, you have to create what is called a multipart e-mail. Much like everything else, this is pretty easy, but is also really easy to break if you are not careful.

The first thing to understand about multipart emails, is that you can nest multipart sections within multipart sections.

If you have ever looked at the International Spectrum Newsletter, there are two sections to it: HTML and Text. That way if you have preference of one over the other, or you are displaying the e-mail on a mobile device, it still displays nicely.

That is how attachments work, which we will also get into later.

The Content-Type

If you remember from the first article, I talked about the e-mail header "Content-Type". This header defines how to process the information within the body of the e-mail. When working with multipart sections, you have to specify the Content-Type as "multipart/alternative" and then supply a boundary id that our information will be in.

Boundaries can be anything you would like, but must be unique within the e-mail. That is right, within that e-mail. You can use the same boundary tag in all your e-mails if you like, but generally the boundary tag is based on the current date, time, and message Id. An example of a boundary tag is:

```
68734-812jasj-o8711n181a
```

Does that mean anything? No. It is a bunch of random characters that are unique within the e-mail. To supply the boundary id, you add it to the end of the "Content-Type".

```
Content-Type: multipart/alternative;
boundary="68734-812jasj-o8711n181a"
```

Now, you define the boundaries sections in your e-mail, but placing a "-" in front the boundary id to start the section, define the Content-Type for that section, add a blank line, and then add your information into the boundary. To start the second section, third section, and so on, you would just repeat the process. Once you are done with all the boundaries, you place the boundary id at the end with "-" on the front and back. See figure 1 for an example.

Gotchas!

There is a big gotcha here. You have to place the HTML and TEXT in the preferred order of display. What this means if you preferred the HTML to be displayed over the TEXT section, then you must place the HTML section as the last in the list. If you prefer the TEXT section to be displayed over the HTML, then that must be the last one listed.

This won't guarantee that the HTML will be displayed over TEXT, since this is a setting on the Recipient's e-mail cli-



**Digital Version
On-Line**

Scan Me!

```
Content-type: multipart/alternative;
    boundary="68734-812jasj-o871lnl81a"
```

```
--68734-812jasj-o871lnl81a
Content-Type: text/plain
```

Hello, this is plain text

```
--68734-812jasj-o871lnl81a
content-Type: text/html
```

```
<html><body>
<p>Hello, this is Html Text</p>
</body></html>
```

```
--68734-812jasj-o871lnl81a--
```

Fig. 1

ent, but if the Recipient has not defined a preferred display type, then your preferred display method will be used.

Now why would you display TEXT over HTML? Wouldn't you want your HTML to always be preferred over the TEXT? Well... mobile phones. While a mobile phone will display HTML e-mail, most HTML emails are designed for a desktop client, not a small two-inch screen. If you are know this is going to a mo-

bile device, then I would recommend place the TEXT as the preferred display, as TEXT will display cleaner on a mobile phone.

Including Different Languages

The nice thing about multipart emails, is you can send your e-mail in different languages, and the recipient can decide which language they want to display

Continues on page 14

Is Your Membership Profile Up to Date?

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

Please go online and update your membership profile today!

intl-spectrum.com/membership



taking multivalue ...
where it has never been before

- Quick and easy to install
- High quality pdf documentation and online help
- Close compatibility with most other multivalue environments
- Maintenance-free file system for ease of use
- QMClient API for development of VB, C and web-based applications
- Very low licensing costs
- No mandatory support contracts
- AccuTerm bundled at no additional cost

Ladybridge Systems Ltd

17b Coldstream Lane, Hardingstone, Northampton, NN4 6DB, England
www.ladybridge.com

US Main Distributor: EasyCo, 220 Stanford Drive, Wallingford PA, 19086 USA
www.easyco.com

www.openqm.com

SENDING E-MAIL FROM MULTIVALUE PROGRAMS - PART 5: CREATING AND SENDING E-MAILS WITH HTML AND TEXT SECTIONS

Continued from page 13

the e-mail in. Again, this is not difficult to do, but it means you have to make sure you provide the information about the languages included in the e-mail.

In order to do this, you have to create a multipart/alternative section to define each of the languages that can be displayed. You will also have to include a header that defines language for each section. See figure 2 for an example.

Content Encoding

When you write your body, you don't have to encode your information in any specific formats unless you want to. By default, the e-mail client will assume the content is in a 7-bit format, which means that all the characters included in the e-mail are between ASCII 0-128.

Well, that is not a problem, correct? That is all the visible characters we care about... except when you have Unicode characters in your e-mail. Then you should define the emails are 8-bit encoding.

To define your content encoding, you add the following e-mail header right after the "content-type" header: "content-encoding: 8bit" There are limitations to 7-bit and 8-bit encoding. Each line of your e-mail can not exceed 998 characters. The likelihood that you run

```
Content-type: multipart/alternative;
    boundary="68734-812jasj-o871lnl81a"
    differences=Content-Language
Content-Language: en, fr, de
--68734-812jasj-o871lnl81a
content-Type: text/plain
Content-Language: fr

Bonjour, ceci est du texte brut

--68734-812jasj-o871lnl81a
content-Type: text/plain
Content-Language: de

Hej, dette er almindelig tekst

--68734-812jasj-o871lnl81a
content-Type: text/plain
Content-Language: en

Hello, this is plain text

--68734-812jasj-o871lnl81a--
```

Fig. 2

```
Encoded:

To illustrate the =
Various rules for encoding,=20the=20following=20=
Example uses a little of

everything

Decode:
To illustrate the Various rules for encoding, the
following Example uses a little of

everything
```

Fig. 3

into this is pretty slim, but it does exist.

Another encoding type you can use is Quoted-Printable. While you don't need it, it can come in handy when you want to provide line wrapping on your TEXT section of your e-mail. The main

issue with TEXT only is that you either provide line breaks yourself, or then hope the e-mail client is displaying that many characters on the screen or you keep your paragraphs all on one line, and hope you don't hit the 998 character limit for 7-bit and 8-bit encoding.

To get around this, you can use the "Quoted-Printable" encoding. The encoding works like the this: No line will exceed 79 characters, any line of text that is part of a line paragraph will include a '=' on the end to identify it as part of the next line of text; ASCII 33-60 AND 62-126 will be left alone, but any other character will be encoded with a '=' and hex value. See figure 3 for an example.

In the next issue we will take a look at addition attachments to your e-mail.

IS

mvpjp.com

MULTI-VALUE PHANTOM JOB PROCESSOR

MVPJP Increases the efficiency of your operations by scheduling and automating jobs to be run in the background. It is cost-effective, user friendly, and highly reliable.





It's All Here!

Stay connected to your MultiValue Community, Technical Resources, and Developer information

- **MultiValue Command Pocket Guide**
- **Technical Resources**
- **User Group Meeting Information**
- **Training**
- **Webinars**
- **News and Articles**



Facebook: intl-spectrum.com/facebook
Twitter: twitter.com/intlspectrum
Mobile: m.intl-spectrum.com

RSS: intl-spectrum.com/RSS
WEB: intl-spectrum.com





PAML Makes Their MultiValue Billing Application 5X Faster with the Caché Database

PAML (Pathology Associates Medical Laboratory) is one of the largest reference laboratories in the United States, serving over 100 hospitals, more than 25,000 clients, and several hundred thousand patients a year. But although there is an ever-increasing demand for their services, PAML had to institute a moratorium on creating new billing services, because their billing systems couldn't keep up. "Our work schedule gives us a four-hour window each night to process the day's accounts," says Sonny Varadan, CIO of PAML, "but the application was taking longer than that to complete the job. We had billing clerks coming in and not being able to start their day because the system was busy."

To address the problem, PAML decided to redesign their system architecture, evolving from a single machine to a more scalable, distributed configuration, and making use of modern technologies such as Web services. That meant they would need to port their billing application from D3 (a MultiValue database) to some other technology. "We

wanted to minimize how much rewriting we would have to do, given the fact that the system had evolved over the last twenty to thirty years," says Varadan, "and we decided that Caché was our best option. Caché supports MultiValue technology, so it provided the easiest migration path. The learning curve for our D3 developers was pretty short."

Over the course of the migration, PAML converted approximately 1,500 routines from D3 to InterSystems' MVBasic, the majority of which were ported without any problems. According to Varadan, the conversion process also highlighted some problems that existed in the original code. "There were some very complicated billing queries that could take hours to execute," he recalls. "In the course of moving to Caché, we were able to identify and fix those problematic queries. Now they run in minutes."

Other performance gains have been similar. Initial results from the newly deployed Caché-based system show that the end-of-day processing which used to take 4-6 hours, now runs in slightly over an hour. The end-of-month processing time has dropped from about 47 hours to 6 hours. As a result, the system is available to all shifts, right from the start of their day. Productivity has increased and PAML has stopped turning away new billing service business. Says Varadan, "I am no longer concerned that the system won't be able to keep up with increased demand."

In addition to enjoying renewed growth, PAML's developers plan

to take advantage of the new capabilities Caché brings to the table. "Caché gives us a lot of flexibility going forward because it can interoperate with a wide range of modern technologies," Varadan says. He reports that PAML's next project is to re-engineer their workflows, using Web services to automate data entry from their clients' systems. And they are investigating using InterSystems DeepSee to embed some reporting functions that are currently handled by a third-party tool.

Varadan is confident that PAML has chosen the right technology and the right technology partner to support growth and new challenges well into the future. He concludes: "Throughout it all, InterSystems has been phenomenally good to us. In my sixteen years working in IT, InterSystems is one of the best software vendors I've ever come across."

For more information, go to www.intersystems.com. ■

wholesale distributors measurably improve their business results.

MITS Distributor Analytics — anchored by the MITS Discover online analytics processing (OLAP) technology platform — provides easy-to-use tools for tracking and optimizing inventory levels and profitability, along with sales trends, cash flow and collections.

"Wholesale distributors use Distribution Management software solutions to run their day-to-day operations," says Fred Owen, MITS President. "These systems generate tremendous amounts of data, but it can be difficult to track key metrics and uncover trends hidden inside them."

MITS Distributor Analytics is designed to extract this complex data from a company's Distribution Management System, then present it to business users in an easy-to-use way. "This lets business managers zero in on the most important aspects of their distribution business—without a lot of tedious work with spreadsheets," says Owen.

Mastering Critical Dimensions of the Distribution Business.

The success of a wholesale distribution business hinges on a surprisingly small set of key business metrics. "It's all about having the right mix of products, deployed to the right places," says Owen. "Match those up with the right sales opportunities, and keep customer payments flowing in, and these businesses can be highly effective in today's competitive environment."

Starting with an executive dash-



New Solution Helps Distributors Boost Sales, Profits, Productivity.

MITS, leading provider of advanced reporting and analytics solutions, today announced a new solution designed to help

board showing KPI metrics that are important across the organization, this new solution provides powerful yet easy-to-use tools for managing each of these business areas. The metrics on this full-screen, graphic dashboard link to underlying detailed reports providing deep interactive analysis. Both the dashboard, and the underlying reports can be set up for automatic distribution to the right users via email, or on demand.

First, the Inventory Return on Assets "Cube" — or analysis-tuned data view — lets business users manage a distribution company's most important asset to maximize its contribution to the bottom line. Users can easily spot weakly performing products and product categories, pinpoint winners and losers, and ensure inventory dollars are invested where they will do the most good.

There's also a Sales Analysis Cube, designed to help distribution business managers quickly identify customers who've stopped buying, sales reps who are lagging behind, and customers who might be draining profits away. It's easy to drill down to — or compare across — branches, products, and other dimensions.

Meanwhile, the Accounts Receivable Cube lets business users quickly spot deadbeat customers, sales reps who depend too heavily on suspect credit, pinpoint unpaid invoices that got lost in the shuffle, and even identify opportunities to improve credit policies.

Finally, the Inventory On Hand Cube helps users investigate

and act on every trend or condition they uncover, by making it easy to see what's in stock and what it's worth. They can also easily identify unusual stock level trends, balance supply with demand, and uncover defective stock and vendor quality issues.

Helping Distributors Win In a Very Competitive Business.

"While inventory, receivables and sales are all critical to distribution success, many organizations overlook another key asset — their information," says Owen. "This new solution unlocks that information, making it easy to understand and analyze."

MITS Discover is a mature, state-of-the-industry analytics solution used by hundreds of organizations worldwide. The new MITS Distributor Analytics is already in use at key customer sites, and is available for licensing now.

About MITS

MITS is a leading provider of Advanced Reporting, Online Analytical Processing (OLAP) and Business Intelligence solutions. The company's product line includes both the MITS Discover OLAP Business Intelligence suite as well as the ad hoc operational reporting solution, MITS Report. Founded in 1996 and headquartered in Seattle, Washington, the company's technologies are used by thousands of organizations, resellers, and systems integrators worldwide. MITS products are firmly entrenched in a wide range of business areas, including manufacturing, distribution, retail sales and services, education, government, healthcare

and insurance as well as in many other industries. ■



Precision Solutions announces the availability of version 4 of the Red Leaf web access portal for the Prelude Advanced

Precision Solutions is pleased to announce the availability of version 4 of the Red Leaf web access portal for the Prelude Advanced Distribution System. This release incorporates an extensive list of improvements including support for product families, ship complete orders, significant product page enhancements, and new corporate user functionality.

Red Leaf offers Prelude customers a unique path to the web by providing a completely customizable user interface wrapped around an extensive web API for integration with the Prelude application. Customers can configure their products for their Red Leaf powered site with more information and options than normally available in the standard Prelude application, with tools to leverage Microsoft Excel and comparable tools for product setup and maintenance.

About Precision Solutions

Precision Solutions (Longmont, Colorado) unravels complex business problems with quality solutions and exceptional customer service. Using our customer centered, results-oriented approach to software development, we're more than just a bunch of technology geeks; we're real people with a passion to deliver solutions that make a real difference.

About Activant Prelude

Activant Prelude fully integrates a distributor's operations. The solution allows immediate transfer of information between departments and permits flexible on-line inquiry into virtually any kind of information. Prelude is a feature rich distribution package with optional integrated modules such as forecasting, integrated customer relationship management, value added manufacturing, and business intelligence. ■



Seinfeld-style Reality Migration Has NGMC Laughing

NorthgateArinso has received another resounding seal of approval after guiding credit information powerhouse National Group Management Corp

Continues on page 18

FROM THE PRESS ROOM

Continued from page 17

through a glitch-free migration from D3 to Reality.

NGMC, a one-stop shop for credit information within the service industry which deals with clients on both a domestic and international level, migrated to Reality last November due to the drastic increase in functionality provided.

Rosalia Letson, Director of Operations at NGMC, admitted that the logistics of the impending switch had been a cause of concern but that the eventual outcome surpassed all expectations, in terms of performance and simplicity.

"It was like an episode of Seinfeld," she said. "Nothing happened, and that was the real beauty of the whole process. It was really incredible, there were zero issues with the whole changeover and the only changes we noticed were the positive ones we had hoped for. We couldn't be happier."

NGMC executives worked hand-in-hand with Northgate's Mark Pick and Mark Fuller, based in Irvine, Calif., plus Tim Gunning, stationed in the company's unit in Belfast, Republic of Ireland, in the four months leading up to the migration.

"From our point of view it is about having a great platform that we know provides the best, most effective and most stable environment for our partners," said Pick, Northgate's Vice President, Reality. "Then we take the time to go through every aspect and ensure a seamless migration for our partner. That is the level of performance we strive for, pride

ourselves on, and which builds lasting and close partnerships with companies like NGMC."

Since the migration the relationship between Northgate and NGMC continues to grow, with the implementation of a joint consulting project that is already providing benefit to NGMC's customer service performance.

About National Group Management

NGMC is a full service industry trade group management organization specializing in the credit, accounts receivable and customer service / supply chain functions. Their industry trade groups are comprised of member companies on a national and international level.

NGMC offers a variety of services in each group ranging from trade experience interchange reports, daily flashes, informative & educational meetings, benchmarking survey's and a great opportunity for members to share information and network. This is all accomplished in an environment that has a strict code of ethics and conduct, adhering to the U S Antitrust Guidelines.

For more information about NGMC please visit www.nationalgroupmgt.com.

About NorthgateArinso

NorthgateArinso (Northgate) serves approximately 1,000 multinational customers, which includes more than 20% of Global Fortune 1000, 4,500 large/ medium customers and approximately 10,500 small/ medium enterprise (SME) customers worldwide. Northgate currently employs over 10,000

employees and operates in 46 countries across 5 continents. The original company was founded in 1969 and since that time has 40 years experience in the IT industry.

Northgate is also one of the market leaders in the MultiValue market where it continues to attract new clients and is considered the industry leader for introducing new technology. Most of the world's largest MultiValue installations across all continents run on a Reality database, many of which are mission critical and make use of Reality's market leading resilience features to ensure maximum availability.

For information about Northgate and Reality please visit: www.northgate-is.com/reality ■

and quickly create reports and charts with a series of mouse clicks. Drilldown scenarios provide users with the ability to slice their data, accessing underlying information on demand. Predictive analytics, flexible result set joining and appending, cross-tabulations and data normalizations together with result set filtering and improved parameter prompting make U2 DataVu 2.0 both powerful and a breeze to use.

Business users will benefit from the new HyperVu web interface allowing them to manage their queries, reports and dashboards and power users will benefit from the new integrated repository allowing easier sharing of any queries, reports and dashboards.

About Rocket Software

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



Rocket U2 announces U2 DataVu Report and Dashboard 2.0 with enhanced data transformation, navigation and 'quick' reports

U2 DataVu 2.0 has been enhanced with a variety of features that make it easier for users to access, analyze and visualize business data. Users can now combine results from queries, apply filters and transformations



Entrinsik Launches Newly Redesigned Website

Entrinsik Inc. (www.entrinsik.com), developer of award-winning web-based operational reporting and analysis solutions for clients across a variety of industries, recently launched a newly redesigned, interactive website. The new site, www.entrinsik.com, provides existing and prospective Informer and SEMtek customers with benefit-

focused content like case studies, white papers, videos and more. The site also highlights a compelling Informer Product Tour that guides visitors through key features and functionality to provide clear insight on how Entrinsik's solutions can meet their specific needs.

Over 1000 organizations around the world choose Entrinsik software for its ease-of-use, flexible architecture, and robust functionality. Entrinsik's Informer Web Reporting is an innovative, award-winning operational reporting and analysis solution used by tens of thousands around the world. Entrinsik's SEMtek ERP software now drives many of the nation's largest professional and continuing education profit center.

"Our website has been redesigned to be more user-friendly, engaging, and informative by allowing visitors to quickly and easily access information about our products and services," says Doug Leupen, CEO and President at Entrinsik. "We are already seeing a dramatic increase in awareness and interest in our solutions."

The redesigned website was developed by Atlantic Business Technologies (<http://www.atlantictbt.com>), a full-service web development company in Raleigh, NC.

About Entrinsik, Inc.

Based in Raleigh, NC since 1984, Entrinsik develops, implements, and supports Information Management Solutions

that enable organizations to maximize performance and improve bottom lines. Entrinsik's Informer Web Reporting is an innovative, award-winning operational reporting and analysis solution used by tens of thousands around the world. Entrinsik's SEMtek ERP software now drives many of the nation's largest professional and continuing education profit centers.

Over 1000 organizations around the world use Entrinsik's software every day. For a demonstration of Informer or a free trial, call us today at 888-703-0016 or visit www.entrinsik.com. ■

eXtra Newsletter

Stay on top of Industry News

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



www.intl-spectrum.com/newsletter

MULTIVALUE COMMUNICATIONS: THE PERSNICKEY PERSISTENCE PROBLEM

Continued from page 7

What's really happening is that the browser is announcing itself to the server on each request and the server is using that announcement to load a specific block of memory for that user from a disk file or database query. When the PHP runtime finishes processing a request, it updates that memory to disk (or database) until the information is needed again. The PHP script may never make another connection to the application. But if it does, all that information is readily available until the script intentionally discards it.

So in contrast to holding a big block of variables in memory indefinitely, PHP holds its information on disk and loads it into memory only when needed, manipulates it as necessary, and then puts it back out to disk before terminat-

ing. This process then repeats until the session is closed by the application, at which point the disk storage is eliminated. (If you think about it, this isn't all that different from what happens to memory when you logoff a Telnet session.)

The disadvantage of extra I/O aside, the advantages to storing state information on disk are numerous. Without a persistent connection there's less chance of a network hiccup causing problems. The actual state information will persist until the application says it's no longer necessary, even between reboots of the server itself! And the memory footprint of the application can be dramatically reduced as memory is committed only when it's needed.

So to answer the question of the necessity of persistent network connections, allow me to offer the oft-quoted words of motivational speaker Anthony Robbins: "If you do what you've always

done, you'll get what you've always gotten".

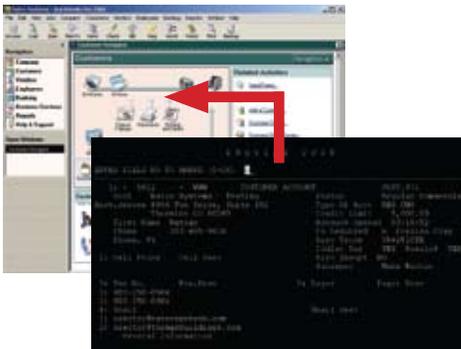
As we look ahead at the next 30 years of MultiValue excellence, are YOU okay with continuing to get what we've always gotten? **IS**



KEVIN KING is the President and Chief Technologist with Precision Solutions, Inc., a leader in technology solutions, support, and training. He is also the author of SB+ Solutions, an enthusiastic private pilot, and Christian guitarist and producer... as time allows.

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



Spectrum University offers once a week training classes in a variety of topics pertaining to the MultiValue market. Looking for training for yourself or employees? Spectrum University Can help! Check out are latest offerings.

Upcoming Webinar Schedule

Syntax Comparison: MultiValue and C#/VB.NET

When working with a new programming language having a syntax comparison can be very valuable. This webinar will show you the syntax differences between MultiValue BASIC and .NET.

Implement ActiveDirectory Security in MultiValue Applications

All MultiValue Applications have security built into them, but it can be a bit disconnected from the rest of the business enterprise. This webinar will show you how to implement ActiveDirectory security within your application.

Back-To-Basics — Free

The Back-To-Basics webinars allow you to review or train new staff on how to use the basics of a MultiValue database. We will be doing classes that cover topics from Dictionaries, Display Formats, Correlatives, I-Types, PROCs, and many others.

MultiValue GUI Applications

After working for years in a Top-Down programming environment, changing to a GUI/Event driven environment can be a challenge. This webinar will cover different aspects of GUI development, include ways to translate "Legacy" applications into GUI environments.

Comparison: MultiValue vs SQL

This webinar will compare the MultiValue database structure to the a SQL Normalize structure. We'll generate SQL tables from MultiValue files, as well as create MultiValue files from SQL Table. These webinars are informative, and are provided to help understand the differences between MultiValue and SQL environments.

Database Triggers

Creating file and database triggers can be tricky with a few gotchas that you need to be aware of. This webinar will cover syntax, templates, and gotchas to watch out for.

For more information or to register please visit <http://www.intl-spectrum.com/webinar>

MultiValue in the Clouds

BY CHARLES BAROUCH

Sadly Turn Away From the Coffee

You know this scenario. You arrived at work, pour a hot cup of coffee, and get ready to tear into that project you keep getting pulled away from and suddenly you hear a co-worker or boss race in and say, “We need to implement <fill in the latest buzzword>. Without <same buzzword> we are just another hopelessly technologically backward company. Fire/Flood/Famine will surely follow if we don’t get on the <same buzzword> bandwagon.”

At this point one of two things happens: A) You sadly turn away from the coffee and spend the next fifteen minutes explaining the buzzword’s real implications to an audience who is unprepared for any answer besides “right away!” B) You sadly turn away from the coffee and spend the next fif-

MultiValue is deployable on the Cloud. Here’s why we aren’t really ready for the Cloud — deployment tools.

teen minutes Googling the buzzword and posting on message boards to figure out what this buzzword means in concept and in workload.

Clouds in My Coffee

Sometimes when we Google and post, it is not because someone else is excited. Occasionally, we allow ourselves to get excited about technology as well. Personally, I’ve been known wax poetic over certain hardware and software from time to time. We are technologists, after all. For some of us, the new twinkling bit of coolness is the Cloud.

We aren’t alone. Amazon, Apple, Microsoft, and lots of other people are screaming, “Cloud!” from every rooftop. So, it wasn’t a surprise when I saw a post on the LinkedIn Reality

group about MultiValue on the Cloud. It wasn’t a surprise that the post said, “Won’t it be cool...”

I leapt in, as did several others, and said, “Can be done, has been done, could be done better.” I thought I’d provide a larger version of my answer for the larger audience here.

Stratus? Cumulus? Cirrus? Cumulonimbus?

A definition is in order. From a delivery standpoint, Cloud is a synonym for Timeshare. I’ve been doing timeshare since the early 80s, and I got into the game late. In essence, Cloud computing is central server computing with shared resources. That sure meets the definition of Timeshare. The difference is in the execution.

From an execution standpoint, the cloud is cluster computing, but at a lower level. Clusters combine several computers to produce the effect of one, faster computer. Cloud deployment allows the combining of memory, disk, etc. to make one computer



Digital Version
On-Line

Scan Me!

beefier, or leaner without reboot, re-configuration, or migration. So, if the execution is done right, a Cloud server is, to the O/S and to the applications, exactly what a non-cloud server is — a collection of hardware.

Cloudy with a Chance...

If a Cloud server running Red Hat Enterprise Linux looks just like a non-cloud server with the same software, MultiValue can run on it. If it looks like Ubuntu and your MultiValue platform can run on that, then MultiValue is deployable on the Cloud.

Here's why we aren't really ready for the Cloud — deployment tools. See, we meet the technical definition of Cloud. But this is a business tech column, so we have to understand how business defines the cloud. Again, it's about deployment tools.

From my original post: "The limit isn't MultiValue, the limit is the 'missing' tools to allow nth number of databases to be hosted separately. Since we have the concept of discreet accounts, the missing part is just a webpage/wizard to deploy you and charge you for the service."

Business loves the Cloud because it offers the following:

1. Less down-time due to migration;
2. Less down-time due to upgrades;
3. Less down-time due to hardware maintenance;
4. Less down-time due to wrestling with software deployment.

Sense a theme here? Business hates down-time. I once had a customer, back in my first IT job, say, "Our system is too important for us to take it down for backups." They literally felt that having no backups was less of a problem than having a down-time event scheduled nightly. Less down-time is a core concept for every business that uses technology, which is to say nearly every business.

If we want to see MultiValue on the Cloud, we can make it happen. There's a lot of money in offering companies a managed system using this sort of technology.

As if by Chance

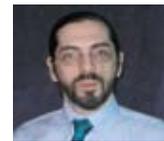
As it happens, I am doing a Cloud deployment right now. While it is MultiValue adjacent — the UniVerse system involved is on a separate server with the Cloud server depends on for content and some layout — the deployment won't require me to move MultiValue to the cloud. Since I know how to install databases to servers, and that hasn't changed, what few things I've learned in this deployment are enough to make me an expert on Cloud deployment for MultiValue. The learning curve is actually that small.

Chances Are

My post oversimplifies, because it only talks about making MultiValue deployable. We all know that MultiValue is a means, not an end. To make us really ready for the Cloud, we need deploy-

ment tools for our utilities (report writers, 4GLs, etc.) , our applications, and for our data.

So, the challenge is out there. In the cloud, people are more concerned with up-time and smooth operations than with brand names. We can stop selling MultiValue and just start selling the productivity that MultiValue produces. Being an embedded database has been a successful model for us, and the Cloud is just the next level of embedding. IS



CHARLES BAROUCH is the CTO for Key Ally, Inc. He is past President of the International U2 Users Group, and a regular Spectrum Magazine contributor.

Contact Charles at Results@KeyAlly.com, or phone (718) 762-3884 ext 1.

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.

Auditing Database Changes with UniVerse Indexing Subroutines

BY CLIFTON OLIVER

The UniVerse database has supported file triggers for some time, not just on UniVerse SQL tables. They are fully compliant with the standard SQL model for triggers — before and after insert, update, and delete events; referential integrity; cascading; the ability to modify or terminate an update; and so forth. Obviously, supporting all this functionality involves a noticeable amount of overhead. But in the Real World, our traditional MultiValue applications manage all of this themselves. The primary use of file triggers I have seen is as a way to capture database change events in order to log them as part of IT governance auditing (SOX, HIPPA, etc.).

In UniVerse release 11.1, Rocket Software has provided us a way of using indexing subroutines to capture changes to a file. To completely understand how this new technique works, let's first review the normal use of indexing subroutines.

The purpose of an indexing subroutine is to derive the index value that is to be cross referenced with the record key. For example, say that we had a need to be able to quickly retrieve all of the purchase orders entered by a particular employee for a given accounting period. Accounting periods do not always follow calendar months. So for our particular industry, in order to determine the accounting period of a particular date, we might have to look up the information in a period table. We could have a subroutine whose job it was to take the purchase order date out of the record, look it up in the period table, and concatenate the period with the user login ID. We then create an I-type descriptor in the dictionary that calls this subroutine then build an index on that I-type. Anytime we add, change, or delete a record, UniVerse is going to want to update that index appropriately.

In the case of adding the record (an "insert" in SQL terminology), the indexing routine calls our subroutine once in order to generate the <accounting period>*<login ID> value that is to have that purchase order number cross referenced to it. Likewise, when deleting a record, our subroutine would be called once to return the index value

from which the purchase order number is to be removed.

Changing a record ("update") requires that our subroutine be called twice — once with the old data to determine which index value is to have the purchase order number removed and once with the new data to determine which index value is to get the purchase order number added to it.

One of the problems with cross-referencing, be it with indexes or the traditional XREF file, is what happens when the majority of records cross reference to the same index record. For example, if you were indexing on an order status field so that you could quickly retrieve all open orders, you would have an index record for the "O" status and another index record for the "C" status code (and, of course, various status codes in between). Assuming that the majority of your orders have been completed, you can see that the index record for the "C" status is going to be humongous. Updating this index record anytime an order goes to "C" status is going to take a lot of time. And it will continue to get worse and worse as it goes along. Eventually you get to the point where a user marks an order



Digital Version
On-Line

Scan Me!

as completed and has time to go get a cup of coffee before they can do anything else with their screen.

What you would want to do in this case is to create your index on an I-type that returns the status code if it was anything other than a "C" and the null string otherwise. Then when you create the index, use the NO.NULLS option. This way, your completed orders (perhaps hundreds of thousands of records) would not be indexed in a single index record.

What does all of this have to do with auditing database changes? Directly, nothing. But once we understand the indexing mechanism and the purpose of an indexing subroutine, we can make use of it for purposes other than for what it was originally intended.

The general idea has been around since the days of Prime INFORMATION. The indexing mechanism is only interested in the value that your subroutine returns. It doesn't care about what your subroutine might do behind the curtain. You create an I-type in the dictionary of the file to be audited. This I descriptor does nothing more than call your subroutine by using the SUBR function.

In the example shown in figure 1, this I-type descriptor will call the AuditSuber subroutine. (The * as the first character indicates that this is a globally catalogued subroutine. If you're not familiar with that concept, see the UniVerse documentation.) This subroutine will return the string value that is to be indexed, which will always be the null string. Now here is where the "trick" comes in. Create the index with the NO.NULLS option and build the (empty) index (fig. 2).

Since your subroutine will never return a value other than the null string, you have essentially used an index as a hook simply to call your subroutine, which will do anything that it wants to do including writing records to other

files, and return to the indexing mechanism which will proceed to do... absolutely nothing.

So what does the subroutine do once it is called? In this case, we wanted to create an audit trail of what has been changed in the data record. The cur-

Continued from page 26

```
cliftonoliv8215 - PuTTY
>ED DICT MY.FILE AUDIT.RECORDS
This is a Type "I" Descriptor last compiled on 08/24/11 at 15:48.
20 lines long.

-----: P10
0001: I
0002: SUBR ("*AuditSubr")
0003:
0004:
0005: 10L
0006: S
0007:
0008:
0009:
0010:
-----: █
```

Fig. 1

```
cliftonoliv8215 - PuTTY
>CREATE INDEX MY.FILE AUDIT.RECORDS NO.NULLS
>BUILD INDEX MY.FILE AUDIT.RECORDS

Locking 'MY.FILE' file for exclusive use.
Starting SSELECT for file 'MY.FILE index AUDIT.RECORDS'.
Compiling "@Ak.0".
SUBR ( *AuditSubr )
Compiling "@INDEX.AUDIT.RECORDS".
SUBR ( *AuditSubr ) ; if @1 NE THEN @1 : ( char ( 251 ) ) : @Ak.ID ELSE

0 record(s) selected to SELECT list #0.

Clearing Index File INDEX.000

Starting DATA processing for index 'AUDIT.RECORDS'!

0 total processed.

Updating INDEX.MAP flags...

Index build of AUDIT.RECORDS complete.
File 'MY.FILE' Unlocked.
```

Fig. 2



Ramp up to full compliance with PRC!

Tracking changes across your software development lifecycle can knock you off course. But you don't have to sacrifice agility to achieve best practices.

PRC gives you the speed, visibility and responsiveness to manage all aspects of software configuration, from service desk through deployment. Productivity leaps when you gain total control.

And we'll be right there, jumping through hoops for you!

SJ+ Systems Associates • info@sjplus.com • http://sjplus.com

AUDITING DATABASE CHANGES WITH UNIVERSE INDEXING SUBROUTINES

Continued from page 25

rent data record is always passed into the subroutine in the @record system variable. So at first glance, it would appear that the job is relatively straightforward. There is, however, one glitch in the technique that has perplexed programmers who use it for years. How do you know why you were in your subroutine? If you have been called because a new record is being created, @record will contain the new data that is about to be written to the file. If you are called because a record is being deleted, @record will contain the data that is about to be deleted. And if a record is being updated, your subroutine will get called twice, once with the old data and once with the new data. So how do you know what @record represents? Old, new, about to be added, or about to be deleted?

- 0 - Not being used from an index
- 1 - INSERT (new) record
- 2 - DELETE record
- 3 - UPDATE derive old index value
- 4 - UPDATE derive new index value

Fig. 3 Values of @IDX.IOTYPE variable

A number of techniques have been tried, usually with stashing things in named common, rereading the record from the disk (so much for efficiency), and trying to guess at what the time interval might be between the first and second call of an update. Frankly, most of these were kludges that were forced to sacrifice either accuracy or efficiency. So there was great rejoicing when true file triggers were introduced.

And then people started running benchmarks. And the rejoicing became somewhat subdued. All of that SQL compliant stuff adds a noticeable amount of performance overhead. So as an alternative, UniVerse 11.1 gives us another system variable to tell us what stage of the indexing mechanism is calling our subroutine. The variable

is @idx.iotype. (Yes. These system variables can be uppercase: @RECORD and @IDX.IOTYPE, if all caps coding is your preference.)

When the indexing mechanism calls your subroutine @idx.iotype is set to an integer value from zero through four. (See figure 3 for the meaning of the five values.) Now when all we want to do is monitor and record database change events, rather than using triggers, we can go back to the lower overhead indexing subroutine technique and know exactly why our subroutine has been called and what we are supposed to do in that instance.

Figure 4 is a listing of a very simplistic auditing subroutine. In real life, you

Continued from page 29

Fig. 4

```
01 *      Subroutine to demonstrate the use of using index subroutines to
02 *      provide auditing of database updates.
03 *
04 *      @idx.iotype = 0 - Not being used from an index;
05 *                  1 - INSERT (new) record;
06 *                  2 - DELETE record;
07 *                  3 - Derive old index value;
08 *                  4 - Derive new index value;
09
10      subroutine AuditSubr(returnValue)
11      common /auditing/ dictAuditTrail, fileAuditTrail, flagOpen, oldDataRec
12
13      equ nil to ""
14      equ otherwise to @true
15      systemMarks = @fm : @vm : @svm
16      printableMarks = "^]\\"
17
18      if not(flagOpen) then
19          open "DICT","AUDIT.TRAIL" to dictAuditTrail else
20              return ; * to caller.
21      end
22      open "AUDIT.TRAIL" to fileAuditTrail else
23          return ; * to caller.
24      end
25      flagOpen = @true
26      end
```

code continues on page 27

Fig. 4 code continued from page 26

```
27
28     begin case
29         case @IDX.IOTYPE = 0
30             null ; * Not being called from an index.
31         case @idx.iotype = 1
32             gosub LogNewRecord
33         case @idx.iotype = 2
34             gosub LogDeleteRecord
35         case @idx.iotype = 3
36             oldDataRec = @record ; * Save before image for next call.
37         case @idx.iotype = 4
38             gosub LogUpdateRecord
39     end case
40
41     returnValue = nil ; * Don't do any indexing.
42     return ; * to caller.
43 *-----
44 LogNewRecord:
45     logEntry = "N"
46     logEntry<2> = @id
47     logEntry<4> = convert(systemMarks, printableMarks, @record)
48     gosub WriteLogEntry
49     return
50 *-----
51 LogDeleteRecord:
52     logEntry = "D"
53     logEntry<2> = @id
54     logEntry<3> = convert(systemMarks, printableMarks, @record)
55     gosub WriteLogEntry
56     return
57 *-----
58 LogUpdateRecord:
59     logEntry = "U"
60     logEntry<2> = @id
61     logEntry<3> = convert(systemMarks, printableMarks, oldDataRec)
62     logEntry<4> = convert(systemMarks, printableMarks, @record)
63     gosub WriteLogEntry
64     return
65 *-----
66 WriteLogEntry:
67     readu dictrecAuditTrail from dictAuditTrail, "&NEXT.AVAILABLE&" else
68         dictrecAuditTrail = "X"
69         dictrecAuditTrail<2> = 1
70     end
71     idAuditTrail = dictrecAuditTrail<2>
72     dictrecAuditTrail<2> += 1
73     write dictrecAuditTrail on dictAuditTrail, "&NEXT.AVAILABLE&"
74
75     recAuditTrail = logEntry
76     recAuditTrail<5> = date()
77     recAuditTrail<6> = time()
78     recAuditTrail<7> = @logname
79     write recAuditTrail on fileAuditTrail, idAuditTrail
80     return
81 *-----
82     end
```



Reality CSV Files

Interacting with delimited file structures for import and export is very important within your enterprise software. Comma Separated Values format is the most common.

Reality has included an easy format for interacting with CSV files. By using the MAKE-SPECIAL TCL command, you can create a Q-POINTER like file that will turn a CSV file into a MultiValue file. This allows you to use standard SELECT/LIST/SORT statements, and MultiValue Basic OPEN/READ/WRITE statements against the data.

You will need to provide a few options to help define how to interact with the CSV files, but the general syntax is

```
MAKE-SPECIAL {filename} CSV {filepath}
```

For example,

```
MAKE-SPECIAL UNION CSV C:\temp\union.csv
```

This will turn the CSV file in figure 1 into a Q-Pointer that displays the output found in figure 2.

The command will create dictionary names based on the heading row. The dictionary names will be created by removing all the extra spaces in the column description and replacing the remaining spaces with periods. Eg, "First Name" becomes First.Name

There are a few options that will come in handy when working with CSV files. Since not all CSV files include headers in the first line, you have the ability to define them yourself.

Include NOHEADING to tell the Q-Pointer that there is no Header line in the CSV file.

```
MAKE-SPECIAL UNION CSV c:\temp\union.csv NOHEADING
```

No Dictionary items will be created when NOHEADING is supplied.

Do you have a Tech Tip to share? E-mail it to editor@intl-spectrum.com

```
"Surname","City","Country"  
"Wilkinson","Newcastle","UK"  
"O'Driscoll","Dublin","Eire"  
"Michalak","Paris","France"
```

Fig. 1

```
:LIST UNION  
Page      1                               13:43:58  06 Apr 2005  
  
Surname... City..... Country...  
Wilkinson Newcastle UK  
O'Driscoll Dublin   Eire  
Michalak   Paris    France  
  
3 Items listed.
```

Fig. 2

FROM THE INSIDE

Continued from page 4

he supplied was information provided on public domain web sites, and a line or two of code on how to access or use the public domain code.

No company process or code was given way.

Issue: Everyone knows how to do this already

Wrong! Wrong! Wrong! I'm sorry. But there are a number of problems and solutions that people think are obvious, but they are not. I see this every year at the Spectrum Conference (Which, by the way, will be April 2 - 5, 2012).

Here is an example: Can you write an article about, "How do you print your MultiValue data to a PDF document?" Response, "Well, everyone knows how to do that!" Did you know how to do it before it was published here in Spectrum magazine?

I don't know about you, but I've found that many people want to do this, and know it can be done, but have no clue where to start.

Just because it seems like common knowledge, it doesn't mean that there aren't people that don't know how to do it.

So, please! I need your help to get the word out to your MultiValue colleagues — MultiValue is alive and well, and it's ready for your mainstream, diverse systems, integrated environments. I want to hear how you solved your business or technical problems.

-NATHAN RECTOR

President, International Spectrum
nathan@intl-spectrum.com

AUDITING DATABASE CHANGES WITH UNIVERSE INDEXING SUBROUTINES

Continued from page 26

will probably want to be doing field level comparisons and more sophisticated recording of who changed what when where.

The named common at line 11 serves two purposes. The first is the fairly typical way of avoiding reopening our audit file every time the subroutine is called. The second is to provide a place to store the "before" record data during the first UPDATE call so that we can compare it with the "after" data provided by @record in the second UPDATE call.

The case structure from lines 28 through 39 is a simple dispatch based on the value of @idx.iotype to the appropriate routine for the type of database change that is occurring. Note that in the case of @idx.iotype type being equal to 3, all we are doing is saving a copy of the old record in named common so that it can be used when we come back to the subroutine with @idx.iotype type equal to 4.

The rest of the code is just to save readable before and after images to our audit file for demonstration and testing purposes. In a Real World subroutine, this is where the actual work of doing record comparisons and building an appropriate audit entry would take place. The results of adding, changing, and then deleting a record with the editor is shown in figure 5.

So does this eliminate the use of regular file triggers? Not really. In cases where

Feedback

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

International Spectrum Magazine has a Feedback Department, sometimes known as Letters to the Editor.

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

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

you need to be able to change the data being saved to the file, be able to allow or disallow the update to continue depending on some set of security criteria, use the other SQL features, triggers are still the way to go. But where your intent is to simply observe what changes are being made to the database and record an audit trail of those facts, indexing subroutines provide a lightweight alternative. **IS**

```
chlfonoliv8215 - PuTTY
SORT AUDIT.TRAIL REPORT 05:42:06pm 05 Sep 2011 PAGE 1
AUDIT.TRAIL Type ID Before After
1 N MY.REC NEW RECORD NEW RECORD
2 U MY.REC NEW RECORD CHANGED RECORD
3 D MY.REC CHANGED RECORD
3 records listed.
>
```

Fig. 5

CLIF NOTES: BLOOD NOT REQUIRED

Continued from page 31

Chicago Manual of Style trying to make up for my perceived lack of formal training. Well guess what? I discovered that not only are there very few hard and fast rules when it comes to grammar and punctuation, but a lot of the “experts” disagree among themselves. And some of the rules that we have had pounded into us, like that chestnut of not ending a sentence with a preposition, are holdovers from grammar books used years and years ago that were not even correct when they were written. They were just some teacher’s personal opinion that got propagated because that particular textbook was popular at the time.

So loosen up and take it easy. Unless you’re in a writing class and doing a paper for some instructor who’s never been published in his or her life, stop worrying about whether or not you would recognize a participle if it was dangled in front of you. Even the pros don’t get it right all the time. I’m

amazed at how many grammar “errors” I find on news web sites like CNN.

Making the physical work easier

No wonder writing seems hard. The way many of us were forced to do it most of our lives is hard, physically hard. Writing essays longhand, re-copying them, having to throw out an entire page if you wanted to switch several of the sentences around? Writer’s cramp is physically painful.

When you get old enough to use a typewriter, assuming your family could afford one at the time, things got noticeably better as far as the physical act of getting words on paper was concerned. But much of the rework when revising your first draft or editing your final copy remained the same. Throw the page out and retype it.

Then came word processors, but they were dedicated devices way out of the price range of most students. It wasn’t until the arrival of the personal computer and word processing programs that the physical act of generating words, rearranging them, modifying

large blocks of text, and producing near-perfect final papers stopped being an exercise in physical torment. Yet I think that most of us still react to “writing” in a negative manner because we are really reacting to the memory of how hard the physical process was.

Write like you were talking to your reader

That’s not a bad bit of advice. But I will go you one better. Don’t write like you were talking to your reader. Actually talk to them. I have been experimenting with dictation software for a number of years. I am happy to report that the technology and the affordable processor power is now readily available to anybody who can afford a personal computer, notebook, or now, even a tablet. In fact, that is how I am “writing” this column. If you want to write articles, if your job requires that you write reports or documentation, or you have to write more than ten e-mails a day, you owe it to yourself to explore this technology.

If you decide to try using dictation software let me share with you a discovery that will dramatically improve your productivity with it. Don’t correct your mistakes while you’re dictating. Every product I have tried has the ability for you to tell it to backup a certain number of characters, delete the previous certain number of words, or correct a phrase for better voice recognition next time. Don’t do it. Pay attention to talking and getting your point across. Then, using Mr. Fingers and Mr. Keyboard, go back through the rough draft and make your corrections and edits the old-fashioned way. It’s much faster than trying to dictate corrections to your dictation. Give it a try. You will find that while you still have to think about what you’re going to say and communicate it in an understandable manner, the actual process of producing the words no longer needs to be a bloodletting experience.

After all, I’ve never heard of anybody getting a “pixel cut.” **IS**

FREE SUBSCRIPTION

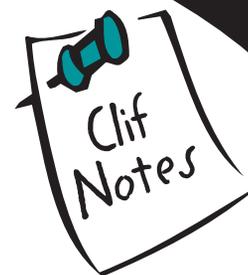
Subscribe on-line at
www.intl-spectrum.com
and keep **Spectrum** magazines
coming **FREE** for 1 year!!





Blood Not Required

BY CLIFTON OLIVER



If you haven't already read Nathan Rector's *From the Inside* column this issue, I would strongly suggest you turn to the front of the magazine and do so. He is requesting help in getting the word out that MultiValue is alive and well and quite able to tackle your modern application requirements. It is also very capable of fitting into a "mainstream" IT shop having a number of different platforms working together. But as long as there are only a handful of us who write articles or blog entries about the kinds of problems we have solved using our various MultiValue platforms and the techniques we used to do it, as a community we are going to continue to appear to be much smaller than we really are. Whereas other platforms have dozens of writers publishing articles for various magazines, journals, and newsletters, in the MultiValue world it seems like you see the same bylines over and over. Not that there is anything wrong with seeing certain writers consistently producing articles! But when those are the only writers you see, it gives the impression that there must not be very many people using this thing, otherwise you would see a lot more people writing about it and discussing it. I think he did an excellent job explaining some of the issues and shooting down some of the myths about what it takes to write an article. Now I would like to add a couple of thoughts then share with you an idea about how to make the process even less daunting.

One of the things that we do each year at the Spectrum Conference is conduct a session for people who would like to consider writing articles for the magazine — what sort of things we're looking for, how to come up with ideas for topics, and how to go about it. So far, these sessions have been very well attended. They seem to generate a lot of enthusiasm, and at the session during the 2011 conference, we were joined by several of our contributing writers who shared their experiences, not only how they wrote their articles, but how it felt personally and how it affected them professionally when their work and byline actually appeared in print. But one of the things that always puzzles me is why with the number of attendees we have, maybe we will get one new writer to do one article. I've given a fair amount of thought to this question and I think I have come up with one of the answers.

Writing is hard

Huh? Really? Especially if you are writing about a topic or technique you know inside and out, have used to solve a real-world problem, and have no trouble explaining to someone sitting across the table at lunch, how hard could it be? And yet, with as many years as I have been writing columns, articles, and technical documentation, I will admit that I sometimes get bogged down even in starting because of the perception that, "This is going to be hard." No-

tice that I didn't say it was, I said that was the perception. And perception is influenced quite a bit by psychology.

Relax, they're just words

Frankly, I blame the school system (public or private). Writing is treated as something that is clearly done either right or wrong. Style is either right or wrong. Grammar is either right or wrong. And punctuation is either absolutely perfect or it is wrong, wrong, wrong, and you are an idiot worth being laughed at. And if you end a sentence with a preposition like the previous one or begin a sentence with a conjunction like this one does, then you have proven yourself to be illiterate.

Nothing could be further from the truth. I am not an English major. I never took an English or writing class that wasn't a mandatory part of the core requirements. My eighth grade English teacher told me, in front of the class of course, that I would never amount to anything because I refused to learn how to diagram sentences. The next year I became the only freshman to have a story published in our high school's literary journal. (So stick it in your ear, Mr. Simmons.)

When I first took the position of Editor of Spectrum magazine, it would take me hours to edit the articles. I spent an extraordinary amount of time with my nose stuck in grammar texts, punctuation guidelines, and

Continues on page 30



INTERNATIONAL
Spectrum[®]
2012 MultiValue Conference

April 2 – 5, 2012

Your Source for Complete MultiValue Solutions

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

International Spectrum 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, Cache, QM, and OpenInsight) software package.

Topics that can be found at the International Spectrum Conference:

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