$7.00 U.S.

INTERNATIONAL Spectrum®

THE MULTIVALUE TECHNOLOGY MAGAZINE | JULY/AUGUST 2014

DATA MINING SOCIAL MEDIA

Asking the Experts

Also in this Issue:

- Responsive Web Design
- UniVerse and UniData Hashed Files – Part 1
- Business Tech: Identity

intl-spectrum.com

April 13 –16 2015

The Premier MultiValue User Event

INTERNATIONAL Spectrum®
2015 MultiValue Conference

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

RETURNING TO THE PGA RESORT, PALM BEACH GARDENS, FLORIDA

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

INTERNATIONAL Spectrum
THE MULTIVALUE TECHNOLOGY MAGAZINE

FEATURES | JULY/AUGUST 2014

DEPARTMENTS

From the Inside

Every year, at this time of year, I start planning for the next Spectrum conference. The 2015 International Spectrum Conference will be April 13th to April 16th in West Palm Beach, FL. That is correct, we are returning to the East Coast for 2015. There are more details in the ad in this issue.

While it seems like the conference is still pretty far away, it really isn't. I have to start to plan out the sessions early, so you have something to sell to your bosses. Part of that planning process is to talk to our regular speakers. We get great content from these industry experts every year, but I also need to know what you as a MultiValue user would like to see.

The demands on your IT department are getting more specific, which leads to specific examples or business solutions that you need to know. I remember those days quite well, back when I had to justify my attendance each year. So I watch the trends that are coming. The trick is finding a way to see which ones are important to the needs of our attendees. Not everything new qualifies.

No matter what I, or our experts decide, I still need to hear from you. The Spectrum Conference is designed for you, the User and Programmer, so 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 talked about virtualization and cloud computing at the conference before. For this year we need to know, are you interested in how to integrate MultiValue databases into a cloud API, like Azure or Google App Engine? Or just how to run a MultiValue server in a Virtual Machine?

Tablets and Cell Phone Apps

What are your company demands: Native Apps, Web Apps, iOS, Android, BlackBerry?

Windows 8/8.1/Metro

If you have not worked with Windows 8, or 8.1, much yet, it can be a pain. Would you like to see general sessions on how to use Windows 8 and Windows Server 2008? These sessions would not likely be MultiValue specific, but sometimes we need to know about other software to make our job easier.

Google Apps

Considering using Google apps in your business? Or for that matter, Office 365? Documents, Calendars, Email, Spreadsheets, all free (Google) or fee (Microsoft) and all can be interconnected to your MultiValue data.

Integration – Email, Salesforce, CRM, EMR, LDAP, Active Directory

Do you have other applications within your environment that you need to integrate with your home grown and vendor applications? We have been an island all to ourselves for so long, we sometimes forget that our business has other systems can use or need MultiValue information for business functions. What integration options do you need?

External Programing Languages

MultiValue BASIC has served us well, but we need to interact with other systems, meaning other programming environments. What languages are you required to work with in your business? .NET, Java, PHP, Python, Javascript, HTML, JSON.... These are just a few examples, but would like to hear from you to see what is most interesting or trends and topics that you would like to have available.

Email me: nathan@intl-spectrum.com

-NATHAN RECTOR

President, International Spectrum

nathan@intl-spectrum.com


twitter.com/intlspectrum


intl-spectrum.com/facebook

INTERNATIONAL Spectrum
JULY/AUGUST 2014

NATHAN RECTOR
President

CHARLES BAROUCH
Editor

SYDNEY BAROUCH
Editor

TRACEY RECTOR
Layout

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

http://www.intl-spectrum.com

MISSION STATEMENT *International Spectrum* magazine's editorial mission is to be the premier independent source of useful information for users, developers, and resellers of MultiValue database management systems, open systems business database solutions, and related hardware, software, and peripherals. Published bimonthly, International Spectrum provides comprehensive coverage of the products, companies, and trends that shape the MultiValue marketplace as well as the computer industry at large — helping its readers get the most out of their business computer systems.

International Spectrum is published six (6) times per year at the subscription price of $40.00 U.S. in the U.S.A.; $45.00 U.S. in Canada and Mexico; $50.00 U.S. for other countries. Single copy rates are $7.00 U.S. in the U.S.A. and Canada, and $9.00 U.S. in all other countries. International Spectrum is published by International Specturm, Inc., 8956 Fox Drive #102, Thornton, CO 80260; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com.

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.

Spectrum University Webinar Series

Upcoming Schedule

Presenting MultiValue data in .NET Grids
09/02/14

Fundamentals:
MultiValue Dictionaries – Week 5
09/04/14

Fundamentals:
MultiValue Dictionaries – Week 6
09/11/14

Integrating SalesForce with your
MultiValue Application – Week 1
09/16/14

Fundamentals:
MultiValue Retrieval Syntax – Week 1
09/18/14

Integrating SalesForce with your
MultiValue Application – Week 2
09/23/14

Fundamentals:
MultiValue Retrieval Syntax – Week 2
09/25/14

Spectrum University offers 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 our latest offerings.

- Weekly Live Sessions
- Only 60 minutes of Your Day
- Expert Presenters from the MultiValue Industry
- Only $175 per connection
- Watch Individually or with a Group
- New Topics Each Week

Have a topic you want covered?
E-mail Nathan with your suggestions at
nathan@intl-spectrum.com

REGISTER AT
www.intl-spectrum.com/webinar

Presented by

SPECTRUM UNIVERSITY

INTERNATIONAL Spectrum

EXPERTS
DATA MINING

Modern businesses rarely have just one database, one programming language, or one reporting tool. We've started a new column called Outside Experts to help us explore non-Multi-Value technologies to see what lessons we can learn. To kick the new column off, we decided to look at data mining, specifically how Nod3x does data mining.

Nod3x: What does it do

Let's look at some of the key features of this product. The interface (figure 1, figure 2) might be bewildering to the new user, but everything has its purpose. The vast array of dots represents every Google+ post which matches the profile #foodanddrink, made since March 28th. We could do the same thing with Twitter, Facebook, or any combination of the three.

If you were running a campaign online and used a hashtag for your marketing, like #is2015, for the International Spectrum 2015 conference, you would see how active the tag was, i.e. how many posts and tweets used that marker.

> ***Lee: Where Google collects website posts, we actually bring together what people are talking about on Google+, Facebook, and Twitter.***

That's the big picture; we can get into the detail by focusing on any dot. We've highlighted one in figure 1. Hovering over will tell you the identity of the person posting it. The size and color of the dot tell you, once you are trained to see it, which are originals, and which are shares and retweets.

Finally, on the right there's a settings menu. Two tabs are particularly useful: Save, which is open on figure 1, and People, which allows you to search for posts and tweets by a specific person. Save lets us pull the data as a CSV, for analysis in other tools. People lets you determine which posts came from your staff vs. the posts representing customers or other interested parties.

Figure 2 is the dashboard view. We get top influencers, the people who promoted that hashtag the most. We also get a word cloud of related tags, where the words are bigger or smaller based on how often they were linked to the target hashtag(s).

Why do we Care?

Just like every other BI dashboard, Nod3x attempts to take massive amounts of data and make it quickly scanable by humans. Just the fact that this is one more tool available to us makes it worth hearing about. However, there's more. Unlike most dashboards, Nod3x is all about un-

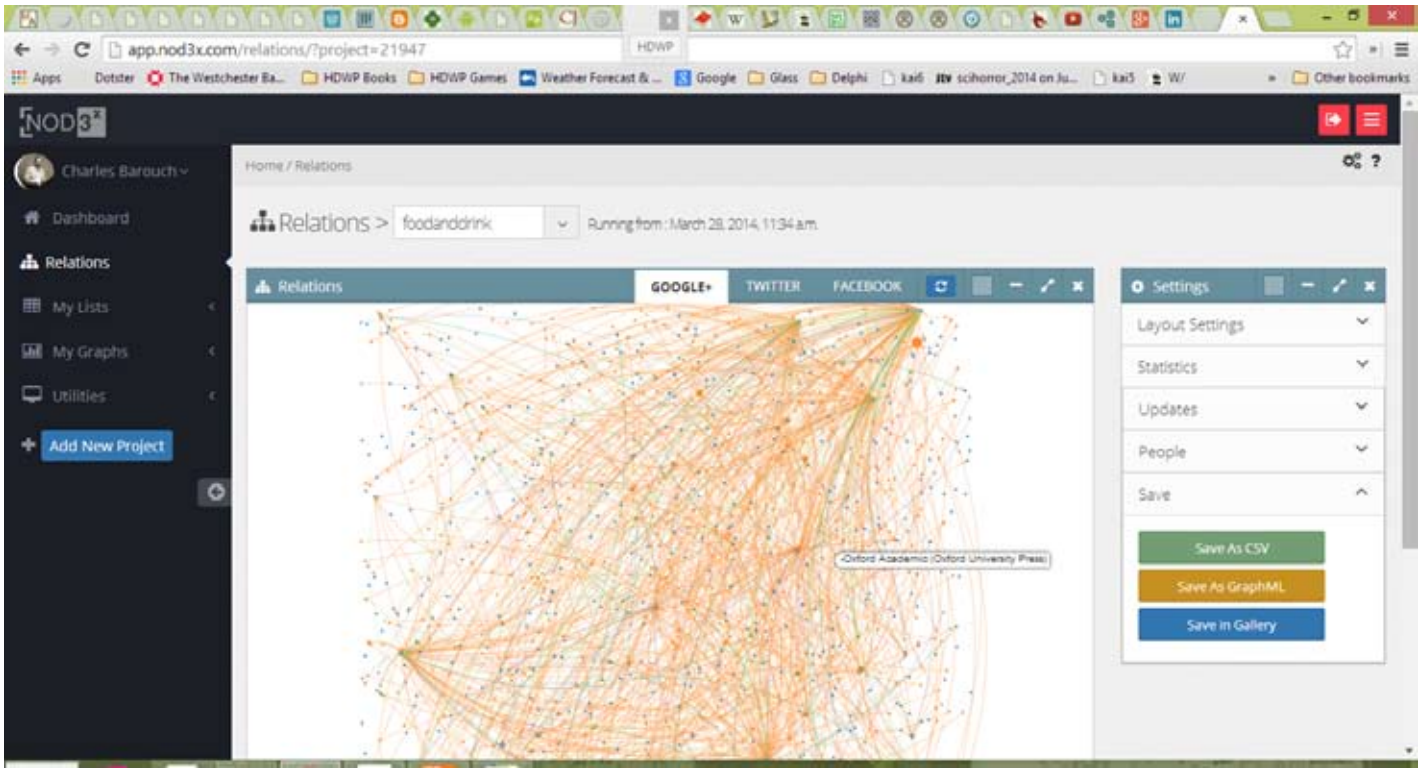

Fig. 1

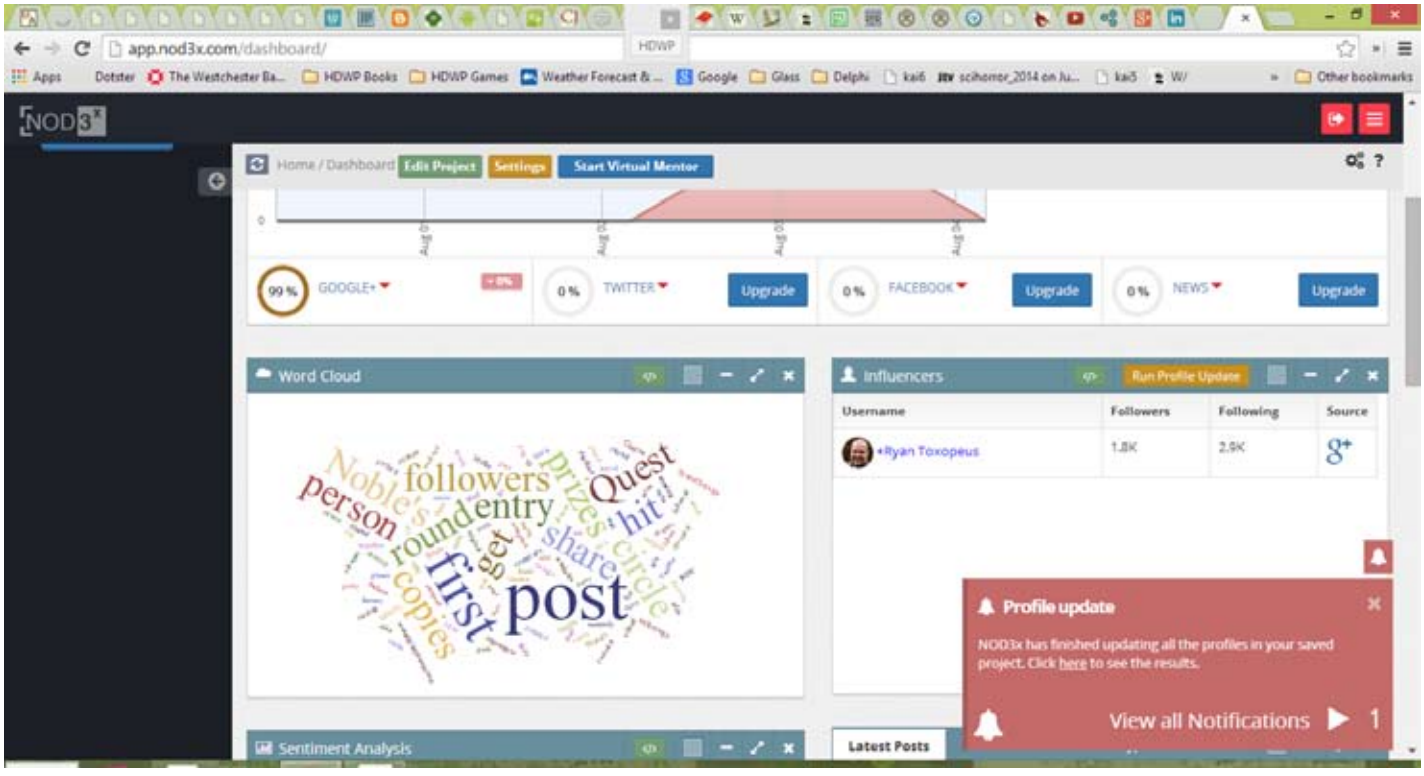

Fig. 2

structured data. As MultiValue technologists, we are constantly working with data deemed 'too messy' for SQL. This makes the ideas behind Nod3x more interesting to us. Lee Smallword and his team are wrestling with the same sort of challenges as we are: how to abstract data and patterns from free-form information.

A good example of this is the Sentiment Analysis tool, which is also part of Nod3x. It attempts to algorithmically quantify posts and tweets into three categories: Positive, Negative, and Neutral.

The Interview

We did a video interview with Lee Smallwood from Nod3x. An excerpt appears below.

LEE: Where Google collects website posts, we actually bring together what people are talking about from Google+, Facebook, Twitter, and others.

IS: You reduce the data into something more visible. Can you talk a bit about that?

LEE: One of the main challenges with social media data is that it is vast. It's volume heavy. It's unstructured. We look at... try to bring in similarities between social media networks: a link, hopefully people are getting the hang of images... a slide show...

IS: You mentioned the graphic element. How important is it?

LEE: Being visual people, we react more quickly to an image than we do to a body of text. As content creators, we have to buy a person's time. That first impact of that image is the key.

The full interview, including Lee's reaction to the ideas behind MultiValue, can be found here: (intl-spectrum/s1067)

Inside Experts

While we interviewed Lee as an expert from outside on data mining, and specifically on social media data mining, we also sent out a survey to our MultiValue software providers and VARs soliciting their input.

Question: What tools do you have or use for data mining MultiValue data?

PICK PROGRAMMER'S SHOP: mv.SSRS when SQL was a requirement. Informer seems to be our goto if the client is strictly MultiValue skillset. Easy to install, and easier to use. However, there is nothing like exposing data to show the client how "bad" their data is (file structures, missing key attributes, mixed use of a single data field) !

KORE TECHNOLOGIES: Most companies use multiple applications within their organization in addition to an ERP system such as: customer relationship management (CRM) systems, eCommerce websites, and other third-party solutions, which in turn generate high volumes of data. To simplify analysis and data mining, this information needs to be consolidated and aggregated into one central location: an enterprise data warehouse.

Products like Kore Technologies' Kourier Integrator Release 4.1.5 can be used to connect and extract data from multiple data sources to build an en-

terprise data warehouse. Kourier consolidates information from UniData, UniVerse and other non-U2 sources (e.g., Oracle, MySQL, Microsoft SQL Server and Microsoft Access) into a Microsoft SQL Server data warehouse in near-real time, enabling more timely analysis and decision making.

Business intelligence and analytics products can then be used to mine and analyze their data. Kore is agnostic to the actual reporting solution, but often recommends its partner's products: CorVu NG from Rocket Software and SSRS (SQL Server Reporting Services) from Microsoft. Both tools are used in the marketplace to make data-driven decisions on a day-to-day basis.

INTERNATIONAL SPECTRUM: Depending on the amount of data that is being processed, we use a combination of SELECTS and REFORMATS along with Excel and Google Prediction APIs to generate reporting solutions.

We have done proof-of-concept on a few other tools, including the 'R' programming environment. Many times MultiValue BASIC subroutines are used with the SELECTS and REFORMAT to generate the raw dataset which is then passed into other reports and analytic engines to get additional results.

Question: How are your customers data mining Social media for use in their enterprise?

KORE TECHNOLOGIES: Unfortunately, the adoption of social media in the MultiValue industry is lagging and we are not yet seeing much use of social media as a primary means of research, education, information sharing and analysis. We are also not seeing much social interaction, which is its real power. We believe increased social media education or training is needed to inform our audience about how to unleash the true power and value that social media offers. Social media is still relatively new however, it's evolving at a rapid rate and we, as a community shouldn't be left behind.

On the other hand, some of our clients are active on social media and sometimes interact with our social profiles, (Twitter, LinkedIn and Facebook). Their data mining strategies are unknown but we do know they use it for customer service, researching company information, learning about product updates, etc. Through other research and communications we've found that in some cases the success (or lack thereof) of a company's social media strategy is unknown to the executives. Many times it's handled by the marketing or public relations person, which leaves everyone else by the wayside. It's interesting because these same people are aware of current marketing and branding initiatives but are left in social darkness.

It's important to note that you should understand your audience before you create a social strategy of how and when to use social media. There are numerous tools, blogs and platforms to use as references and most are free! All in all, it's time for the MultiValue community to overcome the hurdle to embrace the influence and clarity that social media offers.

INTERNATIONAL SPECTRUM: The developers and users that we have talked to use social media data mining to predict mood and find customers that are unhappy. The social media data also helps the enterprise predict what sales may be for specific products based on how much people are talking about them.

Question: What is your future roadmap for data mining and big data analytics?

PICK PROGRAMMER'S SHOP: Continued exploration of tools available. Including those pushing the limits (Intersystems Cache with their "iKnow" technology) of analytics. I'd love to move us toward the difficult data repositories (email as an example) and be able to make analytical/reporting sense of that data.

KORE TECHNOLOGIES: As an industry leader in data management and integration, Kore Technologies is committed to enhancing Kourier Integrator to support the evolving needs for advanced business intelligence, information gathering and data mining requirements that will be needed by the industry. Kourier will continue to extend its architecture to support additional data sources, while continuing to improve the product's ease of use and performance.

INTERNATIONAL SPECTRUM: As more and more open source tools and programs are provided to the code analysis of large data sets, we will provide articles and code samples that will enable MultiValue Databases and LOB/ERP applications to take advantage of the information.

The International Spectrum 2015 conference will also have topics on Big Data and analytics using MultiValue Databases. We will continue to provide MultiValue companies options and solutions using these technologies. IS

"Responsive design uses the technologies that have already been developed. It implements them in a much more integrated way."

Responsive Web Design

BY NEER PATEL

Like with any new buzz word, your boss comes to you and says, "Our site needs to be responsive… We are losing business — make it happen!" And like every buzz word your boss has jumped all over in the past, he doesn't really know what it means, nor how it would affect the business.

What is Responsive Design?

In a 2010 article, Ethan Marcotte coined the term "Responsive Web Design" as a way of describing his approach to creating sites for the ever changing landscape of web browsers. Essentially, this means making one design that would adapt to work on all devices. Ten years ago, this wasn't an issue that needed to be addressed. But now, *everything* has a web browser and users are savvy — and impatient — for all the glitz and glamour. It has to work whether they load it up on their phones, or want to see a full screen video on the refrigerator (It's hard to pinch-zoom when your hands are covered in dough).

The web design communities started integrating techniques and ideas like fluid grid, flexible images, and CSS3 to accomplish this.

Responsive website technology is still a maturing as a design approach, but it is not anything new technologically. Responsive design uses the technologies that have already been developed. It implements them in much more integrated way.

Due to way most websites are currently designed, implementing Responsive web design will require you to totally redo the structure, layout, and likely back end coding. Why you ask? Simplicity. Responsive design makes use of CSS3 classes, DIVs, and CSS Media Queries quite heavily in to create the fluid grid design.

The Power of Media Queries

When it was first introduced, developer jumped all over it because it helped simplify the process of creating pretty layouts. Because the web existed before CSS, many enterprise websites mix CSS stylesheets with inline styles. While this works, it is harder to manage, maintain, and improve.

When CSS2.1 came out, they introduced media queries so that you could create separate stylesheets based on the media: for printing or for screen display. This helped because it caused designers to pull as much inline styles out of the web pages in to the sytlesheets. However, all the layout and positioning of the HTML pages was still done by hard coding positions, or using HTML tables.

While this worked great for full size desktops, it doesn't work for mobile devices. This caused web designers to duplicate their website to make a mobile version. Two code sets to maintain means extra work.

Although CSS3 expanded upon media queries in 2012, developers didn't really stand up and take notice until recently. With these expanded media queries, designers could now define different styles based on the size of the screen:

```
#nav
{
 float: right;
}
#nav ul
{
 list-style:none;
}
@media screen and (min-width: 400px)
and (orientation: portrait)
{
 #nav li
 {
  float: right;
  margin: 0 0 0 .6em;
  border: 1px solid #000000;
```

```
  }
}
@media screen and (min-width: 800px)
{
 #nav
 {
   width:200px;
 }
 #nav li
 {
  float: left;
  margin: 0 0 0 .6em;
  border: none;
 }

}
```

You now have more control over what, where, and how HTML tags are displayed on which device.

Almost Fixed!

This almost fixed the problem of two code sets. Almost, because many enterprise web sites were still used a fixed size layout. The next problem to solve is how to make the layout flexible to present a nice view on a web browser and a nice view on a mobile device.

Enter the Fluid Grid Design. Fluid Grids are a concept designers derived from print designers' Grids. *"A two-dimensional structure made up of a series of intersecting vertical and horizontal axes used to structure content" — Wikipedia.* The axes create a units of predictable fixed sized reference points that are used to design your layout in a consistent manner. The "Fluid" variant still keeps your layout structured, but their size/positions will change depending on the canvas size (or browser window size). Now that is a mouth full.

What is it really? Predefined locations, boxes, and elements which are formatted exactly the way you want them within a predefined space. The difference is that these predefined zones move around on the page instead of staying in a fixed location. The implementation is much like the development of fixed size layouts, just on a ore granular scale.

While text and other elements resize well, images often don't. With CSS3 and media queries, designers can easily alter the size or the actual image itself to create a "Flexible Image"

Making all this work together

Now that you have an overview of what "Responsive Web Design" is, let's apply it to your existing website. A traditional design uses HTML tables or fixed size DIV elements:

```
<table><tr>
<td colspan="2" style="width:600px">
 title
```

```
</td></tr>
<tr><td style="width:300px">
 Block1
</td><td sytle="width:300px">
 Block2
</td>
</tr>
</table>
```

Responsive web design separates zones/boxes (layout elements) that have the information the designer and developer want to display from the actual the layout and positioning:

```
<div class="grid">
 <div class="title">
  title
 </div>
 <div class="block1">
  Block 1
 </div>
 <div class="block2">
  Block 2
 </div>
</div>
```

All a web designer and developer has left in their HTML pages is what is important, not all the extra fluff and chaff needed to position everything.

We use CSS to help position these layout elements based on the layout design and screen size we want. For a small screen such as a smartphone, you would have the layout below:

```
@media only screen and (max-width:
480px)
{
 .grid
 {
  display: -ms-grid;
  margin: 3px;
  -ms-grid-columns: 100%
  -ms-grid-rows: 70px auto auto auto
 }
 .title
 {
  -ms-grid-row: 1;
  -ms-grid-column: 1;
 }
 .block1
 {
  -ms-grid-row: 2;
  -ms-grid-column: 1;
 }
 .block2
 {
  -ms-grid-row: 3;
  -ms-grid-column: 1;
 }
}
```

For a full size screen like a desktop web browser, you would have the layout below:

```
@media only screen and (max-width:
800px)
{
 .grid
 {
  display: -ms-grid;
  margin: 5px;
  -ms-grid-columns: 10px 1fr 10px 1fr;
  -ms-grid-rows: 100px auto;
 }
 .title
 {
  -ms-grid-row: 1;
  -ms-grid-column: 1;
  -ms-grid-column-span; 2;
 }
 .block1
 {
  -ms-grid-row: 2;
  -ms-grid-column: 1;
  background-color: #B2B0B0
 }
 .block2
 {
  -ms-grid-row: 3;
  -ms-grid-column: 1;
  background-color: #EEB0B0
 }
}
```

As you can see, responsive design actually makes web designers and developers job easier. Everything has become simpler.

Gotchas!

Many enterprises still have older browsers in-house. Since Responsive Web design requires CSS3, you have to have a web browser that supports CSS3. So what about all those Internet Explorer browsers that are still below IE version 8?

Some very smart people have already addressed this for you. With a few javascript files, you can implement responsive design on these legacy browsers just like you find on all your newer browsers.

Css3-mediaquieres.js by Wouter van der Gaaf: code.google.com/ Css3-mediaquieres.js/

Response.js: github.com/scottjeh/respond

Fluid Images: unstoppablerobotninja.com/fluid-images/

Conclusion

Unless you really did it right with your enterprise webpages from the beginning, a good 60% of the HTML in those pages is worthless. The key to moving from your fixed size HTML pages to a responsive web design is: figure out what HTML content is important, what should layout elements should be grouped together, and how best to strip the code that used to generate those elements out of your overall web page design.

Once you have the core layout elements isolated, then you have to decide what is going to go where on what screen size. With those answers, you can then build your CSS and media queries to address that set of layout decisions. While this is not hard, once you embrace it, it will be very time consuming.

Using these ideas, you can start building a responsive site from scratch. Or you can capitalize on the generosity of the internet communities and start with one of many frameworks out there. Extra plug for Bootstrap, I find it useful for everything!

getbootstrap.com, getskeleton.com, gumbyframework.com, purecss.io IS

Hashed Files

Part 1

BY PEGGY LONG AND JEFF FITZGERALD

The MultiValue Database and its descendants are known for their flexibility in handling data. Files contain data records that are composed of an arbitrary number of fields, values and sub-values, all of which can be of completely variable length. This is a radical departure from traditional, fixed-length data records found in other SQL databases, and provides a very attractive environment for business applications. The flexibility of the structure places demands on the underlying data storage mechanisms. To meet these demands, MultiValue databases use a hashed file structure. This series of articles will examine in depth the specifics of hashed files as implemented in the UniVerse and UniData environments.

Fixed length records offer a number of advantages from the point of view of data storage. Perhaps the biggest is that the address of any record can be calculated very easily. Similarly, since each data field within a record is fixed length the location of any field in a record can be calculated. A variety of direct access methods can be devised to provide quick data retrieval by depending on these calculations. Think of these traditional database constructs as rows and columns.

Variable length data isn't as easily handled. Since the location of data fields can't be calculated, MultiValue databases use special characters as delimiters — the familiar field or attribute mark, value mark, sub-value mark and segment mark. Locating data is done by scanning the records — for example, the 5th field is between the 4th and the 5th field marks. Since the location of a data record in the file can't be calculated, scanning through a list of records is required to locate a specific record.

One way of limiting the number of records that must be scanned in the search for a specific record is to subdivide the data. If only part of the file must be scanned, it will be quicker to access a record than if the whole file must be scanned. Of course, there must be a mechanism for determining which part of the file holds the desired record. The file storage schemes that we will be examining use a method called "hashing" and thus the files in these environments are called "hashed files". This first article will explore the basic structures and mechanisms involved in UniVerse and UniData hashed files. Note that both UniVerse and UniData dynamic files are "hashed". Later articles will be more specific and look at various aspects of these files in more depth.

Hashed File Basics

The parts into which a file is subdivided are called "groups", because each part contains a group of data records. The number of groups is specified at the time the file is created and is referred to as the "modulo". The word modulo is a reference to "modular arithmetic" which was discovered by the German mathematician Carl Friedrich Gauss in the year 1801.

Modular arithmetic provides a mechanism to sort a file's data records into groups. The data records that share the same remainder when divided by the modulo of the file will all "hash" into the same group. This means that by looking at the remainder we can determine the group the record belongs to and access it by only scanning that group. Less data to scan means faster access!

The Hash Key

Wait a minute! How do you get a remainder for a data record? Well, first each data record has to be identified by a unique record key. The key is chosen by the user or designer of the file and can be nearly anything from a sequential number to a name or a social security number. Since record keys can be non-numeric we have to have a means of converting them to numbers — "remainder" only has meaning when

referring to numbers. The number that the record key is converted to for the hashing process is called the "hash key"

Fortunately, computers already represent characters numerically. There is a code called the "ASCII code" that specifies numeric representations for characters. "ASCII", by the way, stands for American Standard Code for Information Interchange and was developed a long time ago for use with teletypes. As an example, suppose we have a record with the key of "PEGGY". Here are the ASCII codes for the characters in the key:

Character	ASCII Code
P	80
E	69
G	71
G	71
Y	89

Even in environments that use UniCode, an alternative to ASCII which is more friendly to multinational environments, each character is still represented by a number.

Once we've converted the characters to numbers we could invent many different schemes to create a numeric hash key from the record key. In fact, UniVerse offers seventeen different techniques and UniData offers two — these are referred to as "file types" or "hashing algorithms" and we will talk much more about the choices later on. In this article we will use a hashing algorithm that we made up — but it is very similar to the standard hashing algorithm developed for the original PICK MulitValue Database, the UniVerse type 18 algorithm, and the UniData type 0 algorithm.

Here are the rules for our algorithm: Start at the beginning of the record key. Initialize an accumulator to 0. Multiply the accumulator by 10 and add the ASCII value of the next character. Repeat for all the characters in the key. See Here is an example BASIC program that implements this algorithm:

```
ACCUM = 0
INPUT KEY:
*
FOR I = 1 TO LEN(KEY)
   ACCUM = ACCUM * 10
   ACCUM += SEQ(KEY[I,1])
NEXT I
*
CRT
CRT ACCUM
STOP
END
```

So using our hashing algorithm the record key of PEGGY has a hash key of "876899" (Process illustrated below):

Step	Character	ASCII Value	Accumulator
1			0
2	P	80	80
3	E	69	869
4	G	71	8761
5	G	71	87681
6	Y	89	876899

Although the record key must be unique in the file, the hash key will not necessarily be unique. We only need to to calculate the modulo. The remainder will allow us to determine which group to place the record with the key of PEGGY.

Hashing the Record to the Group

Remember that each hashed file has a modulo which specifies the number of groups in the file. The next step in hashing is to divide the hash key by the modulo. The division yields a quotient and a remainder — remember from our discussion of modular arithmetic that the remainder is the important part. The remainder is always going to be a number in the range 0 to the modulo minus 1. If we now add 1 to the remainder it will be between 1 and the modulo — if we number the groups of the file beginning with 1 there is a direct, one-to-one correspondence between the remainder plus 1 and the set of groups in the file. Thus, the remainder plus 1 addresses the group where the data record will be placed.

As an example, suppose a file has a modulo of 7. Using our record with the key of PEGGY, we know that the hash key is 876899. Let's divide the hash key by the modulo:

876899 / 7 = 125271 + (2 / 7)

Quotient = 125271

Remainder = 2

We add 1 to the remainder to get 3. Viola! The record with the key PEGGY "hashes" to group 3 of our file. Whenever our system reads or writes the record, it will know that it resides in group 3.

It is interesting to compare the results from our made-up hashing algorithm with those used by UniVerse and UniData. Create a test file using a modulo of 7 — on UniVerse use type 18; on UniData use type 0 (the default). Now use the RECORD verb to test various record keys and see which group they hash to. Examine the example below of how it should look in Unidata:

```
:CREATE.FILE HASHTEST 7
Create file D_HASHTEST,
modulo/1,blocksize/1024
Hash type = 0
Create file HASHTEST, modulo/7,blocksize/1024
Hash type = 0
Added "@ID", the default record for UniData
to DICT HASHTEST.

:RECORD HASHTEST PEGGY
PEGGY hashed to group 2 and was not found
```

And an example of how it should look in UniVerse:

```
>CREATE.FILE HASHTEST 18 7 1
Creating file "HASHTEST" as Type 18, Modulo
7, Separation 1.
Creating file "D_HASHTEST" as Type 3, Modulo
1, Separation 2.
Added "@ID", the default record for RetrieVe,
to "D_HASHTEST".

>RECORD HASHTEST PEGGY
Record "PEGGY" hashes to group 3 but was not
found.
```

Notice that the group number in the UniData file is 1 less than in the UniVerse file? UniData numbers their groups beginning with 0 instead of 1. So for UniData files, our made-up algorithm could dispense with the final addition of 1 to the remainder. You will find that our algorithm will produce the correct group number for most record keys in UniVerse type 18 and UniData type 0 files. Please don't call us with exceptions that you find! We said MOST because our made-up hashing algorithm doesn't exactly duplicate the real ones.

In our next article, we will take about how the records are stored in "groups", and how this affects read/write and scanning records. IS

Thirty years ago PEGGY LONG and JEFF FITZGERALD created FAST, a utility which can quickly scan a file, report errors and recommend the optimum MODULO and TYPE parameters needed to RESIZE a file.

Their thirty plus years as file maintenance and performance specialists makes Peggy and Jeff uniquely qualified to speak on this subject. Additionally, Jeff brings his experience as a police detective (10 years) and Peggy brings her knowledge as a Ph.D. in Music Education.

Their company, Fitzgerald & Long, Inc. is located in Colorado. If you have questions or comments, please email them at support@fitzlong.com

From the Press Room

Zumasys Acquires AccuTerm Software Leading Cloud Computing Provider Adds to its Growing MultiValue Division with Powerful Remote Access Software

Zumasys, a leading provider of cloud computing and infrastructure solutions to small and medium-sized businesses, announced today the acquisition of California-based AccuSoft Enterprises. AccuSoft is the developer of the AccuTerm software, which provides remote access to MultiValue (Pick) database systems, GUI modernization capabilities and secure access to applications running on the Zumasys cloud platform.

"For decades, Pete Schellenbach and his team have provided impactful client access software to users in North America, Asia, Africa, South America, Europe and Australia," says Zumasys President Paul Giobbi. "We are committed to continuing the strong relationship with AccuTerm's global resellers, independent software vendors and end users, and we are excited about the new places we can take the solution together."

With the intellectual property and technology gained from the acquisition, Zumasys plans to offer new enhancements and services to AccuTerm customers, including:

- New mobile access capabilities for iOS and Android devices
- Zumasys Professional Services to help customers take advantage of AccuTerm GUI's modernization features for character-based applications
- Enhanced mobility for transitioning MultiValue applications into the cloud
- Continued development of AccuTerm technology

With over 30 years of experience in helping nearly 20,000 companies access MultiValue systems, AccuTerm will enhance Zumasys's growing software development team. Pete Schellenbach, founder of AccuSoft Enterprises and creator of AccuTerm, is joining Zumasys as Director of Product Development, working out of AccuSoft's existing Sunland, California, office.

About AccuSoft:

AccuSoft Enterprises was founded in 1980 as a supplier of Pick-based utilities and programs. AccuTerm is the leader in GUI development tools, cloud access and fast, accurate, terminal emulations and connectivity products for the MultiValue market. Learn more at www.zumasys.com/accuterm.

About Zumasys:

Zumasys helps companies of every size transition their infrastructure and applications to the cloud. With Zumasys cloud services, customers can easily access the latest software and hardware technologies over the Web, allowing them to focus on growing their core business instead of managing their IT infrastructures. Zumasys delivers personalized service, integrated disaster recovery and the confidence companies need to outsource the hosting of all their applications, including legacy MultiValue systems. For more information, visit www.zumasys.com. ■

Entrinsik Revamps Corporate Website

With the help of Big Bite Creative, Entrinsik completely revamped their website - www.entrinsik.com - to be easier to use and navigate. New content, information, and testimonials has also been incorporated.

In addition, they have introduced an interactive Informer Dashboard into the site at www2.entrinsik.com/InformerSugarDashboard, enabling visitors to interact with, and get a feel for, how dashboards work. Informer Dashboards has been designed to be fully interactive and customizable, enabling users to monitor and track key metrics and business objectives. Users can quickly and easily build real-time dashboards that integrate data from multiple sources without needing IT to make it happen. ■

InterSystems Zen Mojo Simplifies Development of Web-Based Enterprise Applications for Mobile and Desktop

InterSystems, a global provider of advanced technologies for breakthrough applications, today announced InterSystems Zen Mojo, a new technology for faster and easier development of Web-based applications for mobile and desktop environments. Inter-

Systems Zen Mojo simplifies the development of data-rich, enterprise-class Web applications that leverage the innate functionality and highly responsive user interfaces of mobile devices. The open source plugin framework allows rapid adoption of new and rapidly evolving client libraries.

An evolution of InterSystems' existing Zen technology, Zen Mojo is a lean, efficient framework leveraging JSON (JavaScript Object Notation), the current state-of-the-art format for Web client/server communication. The Zen Mojo technology is embedded in InterSystems Caché®, an extremely fast and massively scalable database, used worldwide in the most demanding Big Data applications.

"Zen Mojo enables us to more rapidly adapt our laboratory management solutions for mobile technologies," said Marcelo Lorencin, President of Shift Consultoria, a Brazilian software company and an InterSystems partner. "Our solutions now provide medical teams with quick access to laboratory test results and electronic patient records from any device. Our strong partnership with InterSystems has helped us increase agility and efficiency for hospitals and healthcare facilities, while also fostering greater patient engagement. Patients can now view their own lab results and, with the right permissions, the results of family members from mobile devices."

"We developed Zen Mojo in response to customer demand for a faster, more efficient way to create cross-platform Web-based applications for both desktop and mobile environments," said Robert Nagle, Vice President, Data Platforms, for InterSystems. "Web application development is evolving very quickly, and the world has moved on from server-generated HTML. The open framework of Zen Mojo allows developers to easily integrate their UI library of choice."

How It Works

Zen Mojo applications are Single-Page Applications (SPA), a popular and efficient approach that reduces back-and-forth communication with the server. Instead of generating HTML that is shipped from server to client, the server transmits metadata necessary for the client to render new content. This metadata drives the layout of the page and includes all the display data as well. The result is a fast, lean approach for developers, while giving users engaging UIs with full access to enterprise data from any device.

Zen Mojo uses a plugin interface for easy integration of state-of-the-art JavaScript user interface (UI) libraries, like jQuery Mobile or the Dojo Toolkit. This approach enables developers to select new client functionality for Zen Mojo as soon as it arises. The technology can also be easily used with frameworks such as PhoneGap for quickly building cross-platform mobile apps.

Zen Mojo is available immediately and can be used with 2013 and 2014 versions of InterSystems Caché. Additional information may be found at www.intersystems.com/our-products/embedded-technologies/zen/ .

About InterSystems

InterSystems develops advanced software technologies that enable breakthroughs. With a passion for excellence and a focus on client success, InterSystems provides data management, strategic interoperability, and analytics platforms used in healthcare, financial services, government, and dozens of other industries. In selected countries, InterSystems also offers unified healthcare applications, based on its core technologies, which deliver on the promise of connected healthcare. Founded in 1978, InterSystems is a privately held company headquartered in Cambridge, Massachusetts (USA), with offices worldwide, and its products are used daily by millions of people in more than 100 countries. For more information, visit www.InterSystems.com. ■

GET CONNECTED.

Knowledge and Education for the MultiValue Professional.

About Our Professional Membership

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

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

Professional Membership Includes:

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

INTERNATIONAL Spectrum®

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

Identity

BY CHARLES BAROUCH

Recently, Syd and I were victims of Identity Theft. In the process of straightening it out — we are still in the process of straightening it out — we've had time to think about how people use and misuse identity.

Who Are You

The primary identifier we all have in common is a name. They are required for birth certificates, driver's licenses, all sorts of documentation. From a tech standpoint, this is a really bad design. I'm sure you've all noticed the lack of uniqueness in names.

In the US, the Social Security Administration tracks baby names. According to the SSA, in 2013, these were the most popular baby names: Noah, Sophia, Liam, Emma, Jacob, and Olivia. You can do your own research on other years here: http://www.ssa.gov/oact/babynames/. What does this mean? It means that if you have a common last name, the odds are that your darling little Liam is going to have a non-unique name.

Pro tip: Lie! I never give a real pet's name or a real mother's maiden name. The problem with lying is that you have to remember which answers you gave to which institutions. That's the other pro tip: don't use the same lies with everyone. Otherwise, one compromise can lead to another.

I'm lucky, in the United States, I have only found one other Charles Barouch. He died in Baltimore, Maryland, in 1928. My father, on the other hand, wasn't even the only person in his family with the same name. His first cousin, who he ended up in business with, had the same first and last. Neither had a middle name.

Despite this, we issue credit cards, security clearance cards, passports… all manner of ID which fully or partially rely on name. As data professionals, what is our obligation toward this type of imprecise information? If our employers have timesheets, non-corporate customers, or mailing lists — just a few examples — we have this sort of data to manage.

Bank On It

Since we had a bank account compromised, let's talk about account security. Typically, a bank handles ID in two key ways: internally we become numbers and externally... we'll get to that in a minute. The bank doesn't see me by name, it sees me as one or more

account numbers. So long as I never interact with the money, this is a perfect arrangement.

Once I want to deposit, withdraw, or check balances, they have to get back to names. The bank can manage my identity in a few ways. First, they have an address for me — either physical or e-mail — and they feel free to send confidential information to that address without any other controls. I might be the five thousandth "Mike J. Smith" doing business with them, but I'm the only "mikey@KeyAlly.com" on record. E-mail address are unique, right?

That works for one-way data. They can't trust the return address on an envelope as proof of identity. They can't trust e-mail completely either. So, when we want to interact, we usually have two options. One is ID card (bank card) and password (PIN). The other is challenge-response.

This is where public information can be dangerous. *What's your first school?* I might be able to look that up. *What's your favorite color?* That might be more secure, although there are only a few common answers.

Pro tip: Lie! I never give a real pet's name or a real mother's maiden name. The problem with lying is that you have to remember which answers you gave to which institutions. That's the other pro tip: don't use the same lies with everyone. Otherwise, one compromise can lead to another.

You are Perfectly You

Are these challenge-response methods perfect? I have personal proof they aren't. Are the ID/password methods perfect? No. What's that leave? Biometrics? Sadly, that's only a little better. Here's the gigantic hole in biometrics: once you scan that eyeball, or fingerprint, or face, once you take that blood sample or skin flake, it becomes data. Quick show of hands: raise your hand if you know how to change data in a computer system. While our bodies are unique, even in twins, triplets, etc., once the data becomes data, it becomes suspect.

What does that leave us? Best efforts.

A Different Angle

Let's take Identity from a different angle, something simpler: deduping. How do we make sure that we capture each unique person exactly once? I tell "Mike J. Smith" from "Mike J. Smith" by looking at their e-mail addresses, for example. Oh, wait. I can't. mikey@KeyAlly.com might also use mikey@intl-spectrum.com for work. Mismatched addresses do not necessarily assure me that I have two different people.

Then we have variant names: Is Peggy Ng also Margaret Chen? They share an e-mail address and Peggy is short for Margaret, but perhaps Peggy is Margaret`s daughter? Score card so far: E-mail, no good. Physical address, no good for the same reasons e-mail is suspect. Real name, not even close. Between nicknames, married names, other legal name changes, duplicate names... It's a mess. IS

CHARLES BAROUCH is the CTO of HDWP, Inc. He can be contacted at www.hdwp.com

MARKETPLACE

ADVERTISER INDEX	
InterSystems, Inc.	Revelation Software
AccuSoft Enterprises	Ladybridge Systems, Ltd.
Kore Technologies	HDWP
PICK Programmer's Shop	Execu-Sys, Ltd.
SJ+ Systems Associates	

HIGH IMPACT 4-COLOR DISPLAY AD

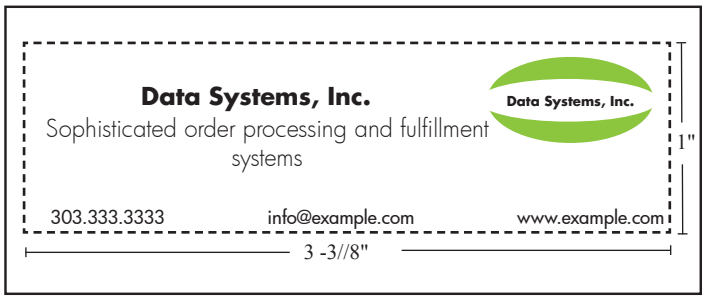

Cost Per Year

1" $1,811.00

2" $3,623.00

BASIC LISTING: $690.00

DataSystems, Inc. http://www.example.com, ... 333.333.3333, ext. 1

- Includes company name, website, and phone number
- Approximately 60 characters per line
- Place under category of your choosing

ADDITIONAL LINES: $280.00 per line

DataSystems, Inc. http://www.example.com, ... 333.333.3333, ext. 1
Sophisticated order processing and fulfilment systems

- Approximately 60 characters per line

ONLINE BANNER ADS:

VERTICAL BANNER (120X240)

$2,880.00 per year or $300.00 per month

RECTANGLE (180X150)

$2,592.00 per year or $270.00 per month

SQUARE BUTTON (125X125)

$1,680.00 per year or $175.00 per month

E-MAIL NEWSLETTER BANNER:

HALF BANNER (300X90)

$4,200.00 per year or $175.00 per issue

ONLINE MARKETPLACE LIST:

BASIC LISTING – FREE

- Company Page
- Product Page – Limit 1
- ADs will be displayed on your page

PROFESSIONAL LISTING

$350.00 per year/ $35.00 per month

- Company Page
- Product Pages - Unlimited
- No ADs displayed on your page
- Download Links

For more information contact Nathan at: nathan@intl-spectrum.com

MARKETPLACE

ACCOUNTING

Natec Systems
www.natecsystems.com | nrector@natecsystems.com

COMPLIANCE

SJ+ Systems Associates
www.sjplus.com | sjoslyn@sjplus.com

CONSULTING

Clifton Oliver & Associates
www.oliver.com | wco@oliver.com

Drexel Management Service
www.drexelmgt.com | dconboy@drexelmgt.com

Execu-Sys, LTD
www.eslny.com | mh@eslny.com

HDWP
www.HDWP.com | results@HDWP.com

PICK Programmers Shop
www.pickprogram.com | brian@pickprogram.com

Precision Solutions
www.precisonline.com | Kevin@PrecisOnline.com

CREDIT CARD PROCESSING

MV Tech Services Inc.
http://www.mv-tech.net | tom@mv-tech.net | 952.474.3795
Process credit cards directly from any MultiValue application.

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

AccuSoft Enterprises
www.asent.com | pjs@asent.com

WANT TO SEE A SPECIFIC TOPIC?

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

E-mail: nathan@intl-spectrum.com

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

International Spectrum Magazine

on Paper, Online, and Beyond

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

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

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

INTERNATIONAL Spectrum

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