INSIDE: Developing User Interface Standards

INTERNATIONAL
Spectrum®

THE MULTIVALUE TECHNOLOGY MAGAZINE | MAR/APR 2011

PLUS!
Who Owns the Data?

Character Encoding

Getting Ready for the World Stage

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 software technologies for breakthrough applications.

InterSystems

INTERNATIONAL Spectrum
THE MULTIVALUE TECHNOLOGY MAGAZINE

FEATURES | MARCH/APRIL 2011

DEPARTMENTS

From the Inside

BY NATHAN RECTOR

Welcome to the International Spectrum 2011 Conference!

For those who are here in Florida, Welcome, Enjoy... I'm sure we will be talking throughout the conference.

For those that were unable to make it, don't feel left out. Spectrum 2011 has two events: Live and Virtual. Our Live event is full of education and networking, I know many of you wanted to join us but were unable to join us due to budget constraints.

This is where the Virtual conference comes in. April 12–14 are the dates for the Virtual part of our conference. Now, the question you likely have is — what is the differences?

Well, it's not different but an extension of the Live conference. Trying to find all the solutions available to you can sometimes be hard and time consuming, but the Live and Virtual conference allow you to see all the products and services in one place.

The Virtual Conference provides you access to our sponsors to either initiate new conversations, or continue ones that you have started at another time. Each of our sponsors has access to our Instant Message servers to schedule available times when they will be free to talk with you. Some of the sponsors may be providing you the ability to talk with different people within the company as well.

In addition to the Instant Message chats, you can access and download On-Demand Videos provided by the exhibitors. These give you even more information about vendor products and solutions.

Some, though not all, of the Live conference sessions will also be available On-Demand during that time as well. The On-Demand session will give you the basics you may need or want to get started with a solution to whatever problem you are looking to address.

For those of you here at the Live conference, you get access to the Virtual content as well as part of your attendance to the Live Conference. Also keep a look out around the conference for the QR code hunt. You may win something.

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

twitter.com/intlspectrum   intl-spectrum.com/facebook

Application Developers just got a brand new bag.

And we continue to fill it with the tools you'll need to get the job done; this time with SQL Connector and Source Code Management.

SQL Connector provides the ability to easily connect OpenInsight to a SQL database. With this connector, SQL tables are now seen as native tables to all components of OpenInsight including forms, reports, popups and stored procedures. The O4W web development toolkit can also be used against SQL data once the SQL data is attached to OpenInsight. A very handy tool, indeed.

Source Code Management allows developers to save a unique copy of every program compiled, as it is compiled, and allows for easy versioning of those compiled programs. Source code can be grouped in modules and the different versions of source code can be compared to evaluate differences.

You won't believe how easy it is to make your applications and data available anywhere when you develop them with OpenInsight 9.2.1.

Visit our website and download your evaluation copy today.

Character Encodings

B Y   B R I A N   L E A C H

Historically, few MultiValue developers have needed to immerse themselves in the murky waters of character encodings. But in today's world, with ever greater internationalization, a basic grasp of encoding is becoming essential — whether working with email, web pages, client languages such as those in the .NET platform, or for data storage and transfer.

Unicode and internationalization are far too large a subject to cover in a single article, so this will give a high level understanding of how modern character encoding works, how it came about — and why you need to know about it.

The ASCII Era

MultiValue databases typically use an eight bit encoding scheme, which dates back to the days when all computers spoke English and IBM was still clinging on to EBCDIC whilst the rest of the industry was standardising around ASCII. Within the ASCII model all necessary characters in the English speaking world could be fitted neatly into seven bytes along with certain control characters and punctuation. This even left a whole bit to spare on the then-prevalent eight bit architectures, and different technology groups decided to use this for their own devious purposes: giving rise to our own marker characters — the field, value, subvalue and text marks — sitting at the top of the table.

ANSI

ASCII was fine and dandy so long as you spoke English, but inevitably those pesky foreigners wanted in on the act and all those extra characters from 128-255 proved an irresistible lure. The ensuing free-for-all created complete chaos until the ANSI standard came along to put things right. Under this standard, the first 127 characters were kept for good ol' ASCII (required for programmers — since all major programming languages and operating systems are still derived from English today) and those from 128 up assigned to various regional languages through the application of different code pages. So by setting your Windows PC to use page 874 you could render Thai characters or to page 1251 if you wanted Cyrillic. Similarly if you look through a VT100 emulation manual, for example, you will find selection commands between US English, Finnish, and French code pages, amongst others.

All of which was good news, so long as you only ever worked within one region, and you didn't need anything too complex like many of the far eastern character sets. But it also set an important precedent.

Below 128, all characters were predictable, and so a character and its encoding went hand in hand. If you are an English speaker, you can treat the two as synonymous: CHAR(65) is a capital A is a capital A is always a capital A. Above 127, and a character *only* makes sense if you understand its context, which in the case of the ANSI standard meant knowing the code page for which it was intended.

This presents some problems. Receive character 138 as part of a document from an international partner and you may need to know whether to display a Cyrillic capital letter LJE , a capital S with a caron or any of a dozen different variants. In short, knowing a character value alone is not enough. The link between the character value and its intended meaning had been broken.

Unicode

The ANSI code page approach had three huge drawbacks: it was only useful if your character set could be crammed into the extra 128 characters available in a code page, which ruled out a lot of the far eastern regions; it only really worked if your language was composed of discrete characters, again ruling out languages that use combining characters to form a single symbol; and there was no standard way to com-

municate which code page to use. Oh and naturally enough, DOS, Windows, and various types of UNIX all tended to use different code pages.

Clearly what was needed was one single standard to unite them all into one series — one that could define every possible regional character, symbol, combining character and, for good measure, various mathematical and engineering symbols, while not forgetting historical languages like Anglo-Saxon, Ogham , hieroglyphs, and mythical languages like Tolkein's Tengwar and Cirth(fig. 1).

So the computer industry did what it does best and bravely set up two competing standards: ISO 10646 defining the UCS (Universal Character Set) and Unicode. Fortunately for all, reason ultimately prevailed and the two standards harmonised, so for developers today Unicode is the standard of choice.

But that is far from the end of the story.

A common misconception is that Unicode operates similarly to ASCII but with a lot more characters, and that these are represented as double byte (16 bit)...

Unicode and Character Encodings

A common misconception is that Unicode operates similarly to ASCII but with a lot more characters, and that these are represented as double byte (16 bit) giving a maximum of 65,536 possible entries.

One reason for this misconception is that Windows originally offered double byte characters as a way of encoding far eastern sets, and development tools such as Delphi and C# currently define a 16 bit char data type. If you are an English or European language speaker, save a Notepad document as Unicode and it will indeed write the content as a series of double byte characters with the unused byte set to zero.

True double byte has the advantage of consistency. It is just as easy to move forward and back through a regularly double byte string as it is through a single byte string, so long as you can ignore brain-dead C routines that expect 0x0 as a string terminator and you know the order in which the two bytes are combined. The question of big-endian or little-endian is something over which manufacturers cannot agree, so your Notepad document also includes a lead-in character known as a Byte Order Mark or BOM which shows as 0xFEFF or 0xFFFE depending on the architecture and which can be used to identify which way round these are coded. Character handling routines are expected to byte swap as required. But Unicode is <u>not</u> a 16 bit encoding standard.

In reality, Unicode and UCS are nothing more than classifications. Their purpose is to simply define a global code table that assigns a positional value to each character (known as a "code

Fig. 1

	E00	E09	E0A	E0B	E0C	E0D	E0E	E0F
0	ᛈ	ᚌ	ᚐ	ᛞ	ᛗ	ᚺ	·	
1	ᚱ	ᛚ	ᚌ	ᛈ	ᚼ	ᚦ	⁞	
2	ᚴ	ᛍ	ᚏ	ᛗ	ᚺ	ᚴ	⁞	
3	ᚨ	ᛁ	ᚋ	ᛊ	ᚺ	ᚼ	⁞	
4	ᚼ	ᚨ	ᚢ	ᛁ	ᛗ	ᚦ	‖	
5	ᛒ	ᛆ	ᛏ	ᚺ	ᛉ	ᚦ	̂	
6	ᚨ	ᚴ	ᚴ	ᛁ	ᚼ	ᛁ	̲	
7	ᚦ	ᛎ	ᚴ	ᚴ	ᛘ	ᛁ	̥	
8	ᛔ	ᛏ	ᚴ	ᛉ	ᚼ	ᛌ		
9	ᛂ	ᚤ	ᚵ	ᚵ	ᚢ	ᛌ		
A	ᛤ	ᚤ	ᛏ	ᚼ	ᚹ	ᚴ		
B	ᚱ	ᚤ	ᛡ	ᚼ	ᚢ	ᚦ		
C	ᛈ	ᛚ	ᚷ	ᚼ	ᚨ	◇		
D	ᛐ	ᛚ	ᚼ	ᚼ	ᚢ	ᛞ		
E	ᚤ	ᛁ	ᚲ	ᚻ	ᛉ	ᚢ		
F	ᛏ	ᛐ	ᚷ	ᚺ	ᚤ	ᚹ		

How UTF-8 Encoding Works

UTF-8 uses a progressive scheme to encode characters, in which the most popular (to the American audience) characters are held in a single byte and the more obscure (to the same audience) expand into up to 6 bytes.

UCS characters U+0000 to U+007F (ASCII) are encoded simply as bytes 0x00 to 0x7F (ASCII compatibility). This means that files and strings which contain only 7-bit ASCII characters have the same encoding under both ASCII and UTF-8.

All UCS characters above U+007F are encoded as a sequence of several bytes, each of which has the most significant bit set. Therefore, no ASCII byte (0x00-0x7F) can appear as part of any other character and all possible 231 UCS codes can be encoded.

The first byte of a multi byte sequence is always in the range 0xC0 to 0xFD and it indicates how many bytes follow for this character.

The bytes 0xFE and 0xFF are never used in the UTF-8 encoding, which is good news for field marks.

The following byte sequences are used to represent a character. The sequence to be used depends on the Unicode number of the character:

U-00000000 – U-0000007F: 0xxxxxxx
U-00000080 – U-000007FF: 110xxxxx 10xxxxxx
U-00000800 – U-00007FFF: 1110xxxx 10xxxxxx 10xxxxxx
(etc)

CHARACTER ENCODINGS

point"). What they specifically **do not** define, is how those characters should be stored or encoded. And so far, they define a 32-bit character set (2^{16}), so double-byte encoding just isn't enough to hold all possible entries.

Now if you're an English speaker still working with the 7 ASCII bits, the thought of all those extra bits might be a scary and you may feel aggrieved at having to store them all. And then there are all the legacy data and documents that predate Unicode — what should you do with them?

In fact, there are many different ways in which Unicode characters can be encoded: BMP, UCS-2, UCS-4, UTF-7, UTF-8, UTF-16 and UTF-32 amongst others, along with the old code pages that can show some, but not all, characters. But out of this bewildering series, the chances are you will be introduced to either UTF-16 or UTF-8.

UTF-16 is the regularly spaced double byte encoding we just met. That's good enough for most international text unless you really have a burning need to display Elvish or some mathematical symbols. It is not compatible with ASCII but it is relatively easy to navigate and if your data is in ASCII you can easily translate between the two by either adding or stripping off the zero byte. But even so, that's extra work and who needs that?

So the popular alternative in the Western world is UTF-8, the brainchild of Ken Thompson who famously devised it on a New Jersey diner placemat.

UTF-8 uses a variable number of bytes to encode a character — anything between one and six bytes. This makes it expensive for documents that use characters at the far end of the Unicode table, but has the huge advantage that the first 127 characters can fit into a single ASCII-compatible byte. So you can take your ASCII document, mark it as UTF-8 and sit back knowing that you are now in the modern world.

Specifying the Content Type

Handling text is no longer just a matter of scanning bytes. In today's world knowing the encoding scheme matters. For email, this is specified using the Content-Type header. For a web page, this uses the meta tag:

```
<meta http-equiv="Content-
Type" content="text/html;
charset=utf-8">
```

Note that this is actually *inside* the HTML document and the browser needs to parse past the <html> and <head> tags to get that far.

Which just proves that text encoding isn't something to be scared of but it is something that needs to be understood. **IS**

BRIAN LEACH is an independent MultiValue consultant living in the UK. He is current President of the U2 User Group and author of a number of books on UniVerse programming.

User Ownership of Data

BY CHARLES BAROUCH

People Who Live in Glass Houses Shouldn't

In the beginning, data belonged to the hierarchy of priests who dwelt in the rooms of glass, where cleanliness was more important than godliness. And the priests spoke unto the people, and they said, the data is ours and you shall not partake of it without a proper help desk ticket. You shall not partake of it until our schedule allows. Our claim is greater than yours. You do the mundane work of generating revenue, but we spend our days in the contemplation of the code.

And the people were demoted from co-workers to mere users. And the priests looked down upon them and thought ill of their priorities. It was a time of great darkness and, eventually, the very fabric of computing was rended and remade until the priesthood fell from grace.

A Byte in The Hand

Computers moved out of the clean rooms. They were forced to move out. Frustrated users brought PCs and other tools into the workplace to allow them to directly manipulate their data, their way, on their schedule. The spreadsheet became a weapon in the war to get around the priests of IT. Even those shops where the glass houses still ex-ist have seen the sea change. Workers want to get work done, and if IT is perceived as a bottleneck, there are ways around them. Remember, bottlenecks create frustration and frustration can turn to anger.

Unfortunately, some of that anger has been turned against IT people who never acted like high priests. Oddly, some of the people who have taken up this negative attitude weren't in the workforce when the war ended. Some weren't even born at that point. It's an exercise in conditioning that would make B. F. Skinner proud.

For better or worse, the data *is* in the hands of the users. Smart IT people see these people as co-workers, as creators, and owners, not just as "users." All in all, it is for the better. So, if user's own the data *and* user ownership is a good thing, *and* we like to advance the cause of good things happening, what does user ownership of data mean to the IT professional?

Look Before You Leak

One thing it means is that Security is much harder. Just like in every other part of life, when something goes from "tightly controlled by a handful of people who operate in a single group" to "generally available to a wide variety of people with diverse skill levels and training" the ability to manage access becomes massively harder. For an easy example, in the glass house days, security had little to do with user rights management. Now many organizations depend on one or more full-time LDAP administrators to control rights. When a task which used to be a minor project becomes a full time job, that's a sign of a major change in the rules.

Even if you manage security to five nines precision, you can still have employees who do not understand the sensitivity of the data they are allowed to access. If Sally in accounting talks about that million dollar banking snafu at a cocktail party and Martin, who works at a competitor, overhears... So, when data moves from the few to the many, we as good citizens of our company need to raise the issue of staff confidentiality training.

Of course, inappropriate dissemination of data is often less dangerous than misunderstanding data.

Baa, Baa, Spread Sheet, Have You Any Wool

Joe plugs in his data, collected from the corporate system of record, and runs his many spreadsheet formulas against it. He proudly sends the results up to his manager, Angela. She looks at his

work and concludes that the corporate system needs to be replaced because this is the seventh time this month that she has gotten conflicting data. Joe's report shows a gain of 5% over last year. Frank's report shows a loss of 3% over the same period.

Why? Is it because one of them did a bad job in manipulating the data? Maybe. It could also be that Joe has more access to data than Frank, so Frank's numbers only include three divisions while Joe is including five divisions. It could be as simple as "What does 'day' mean?" Joe might be including weekend sales into Monday, which at the break of the year pushes a lot of money from December 30th (last year) to first week in January (this year) which changes both totals. It could be because Joe is looking at Dollars and assuming gross, when Frank is looking at the same column and knows that it is net of standard overhead. Perhaps Frank pulled his data three months earlier.

When you have a limited number of people doing the reporting, these

> *For better or worse, the data is in the hands of the users. Smart IT people see these people as co-workers, as creators, and owners, not just as "users."*

things are easier to sort out. Reducing the number of people does not eliminate the errors, but it does allow for rehabilitating the errors faster.

Data, Data, Quite Contrary

So, user ownership of data is a bad thing, the priests were right, and everything should go back to how it was. No. No. No. User ownership of data is more complex, but it is a good thing. To start with, IT knows if the data ties out but users know if the data is correct. For example, when I test order entry, I'll set up a customer and buy some products. When the sales people

test order entry, they will expect it to reject certain combinations of products because they can't be shipped together. They'll expect that our fireworks division shouldn't be allowed to ship into certain states. They'll expect that quota will be checked on exports from China to the US on some classes of goods.

The problem with IT ownership of data is that we live on the wrong side of the screen. People who live and die on the details of a particular transaction set are better able to see the holes. They can tell if we've made things tight enough. They know if it is too restrictive. IT needs to partner with the users. We need to elevate each other to the level of co-workers.

What we bring to the table is knowledge of the architecture of the data. We know how it glues together. We know if a particular piece is being collected. We know where it is being used. Just as we know the structure (schema), the users know the content and the

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

BY NATHAN RECTOR
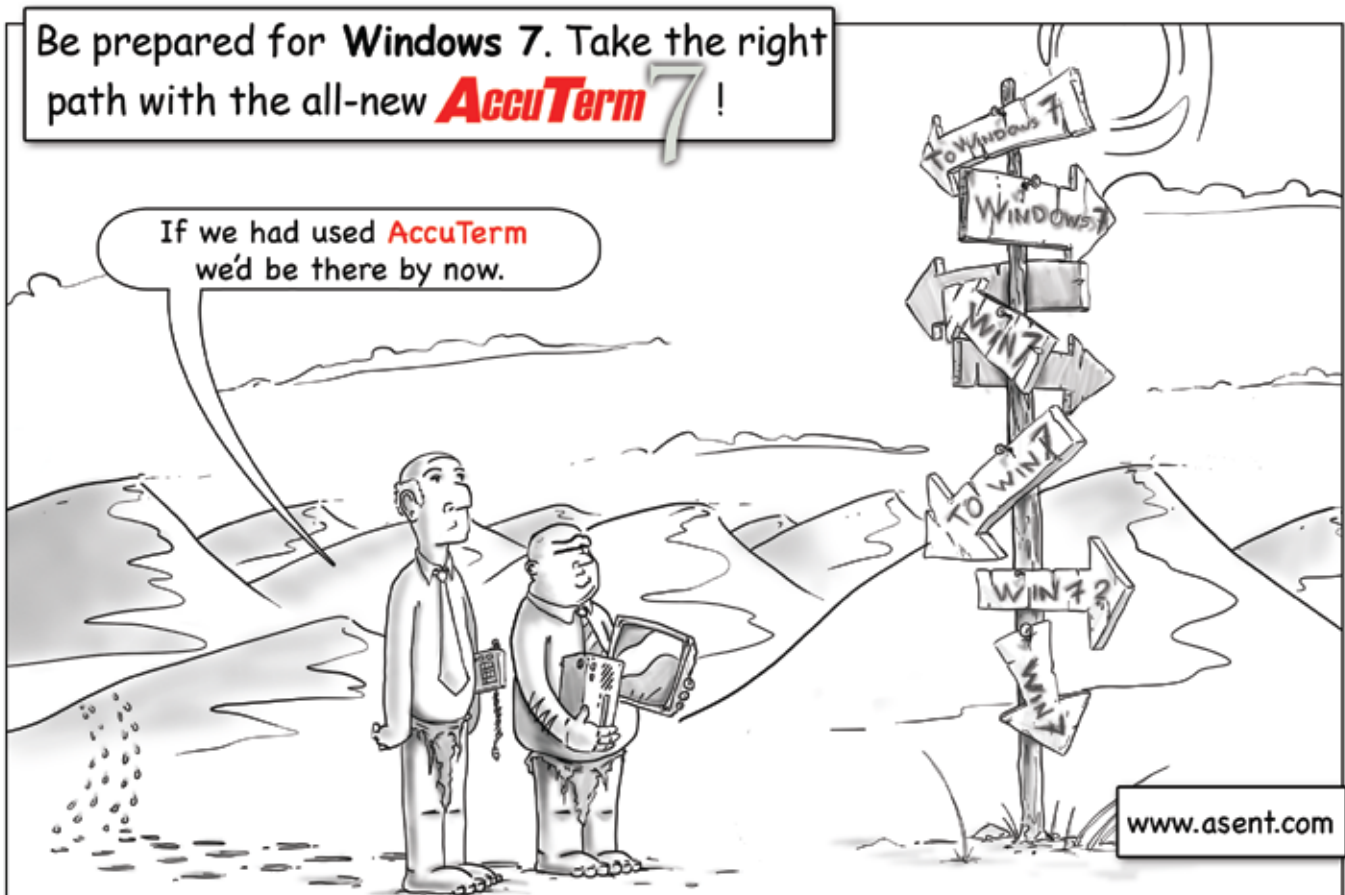
Sending E-mail
from Your
MultiValue Programs
Part 2 – Using IIS

L ast issue I talked about how to create a simple e-mail, with headers and body and all. In this article I'll start talking about how to take the e-mail you formatted and get it sent from your system.

You have several different options to consider, depending on which host system your MultiValue database is on.

- IIS (Windows)

- SendMail (Linux/AIX/*nix)

- Blat (Windows)

- Exchange Email Server (Windows)

- Native SMTP client

There are many other options depending on the SMTP server you use in-house. Some of these have easy e-mail sending interfaces. Others require some scripting, but all are available to you.

While sending e-mail using an SMTP server is easy, getting your SMTP server setup and accessing the Internet correctly is a much larger topic. When setting up an SMTP server or client, you need to consider all the security features the SMTP server may be implementing.

Spam filters, Relay authentication, Domain Keys, and many other issues can cause your e-mails to be rejected by the server. This may look to your users or systems admins as something being wrong with your e-mail sending program, when in reality it is outside of your control.

An example of this is when your company e-mail admins implement "Relay Authentication." This requires you to provide username/password along with your e-mail message to your sending program. This can be problematic if you are using TLS encryption for relaying.

While these are important issues, they are outside of what we are going to cover in these articles.

Since we have to start somewhere, we will start with IIS (Internet Information Server).

Why IIS?

Most of us have a Windows XP, 2000, 2003, or 2008 server on site. This means we have an IIS server available to us to send e-mail with. IIS includes a SMTP service that you can integrate into your database process without any extra third-party tools or command-line tools.

The only requirement is getting the e-mail written to the Windows machine that you have the SMTP Service installed on. Writing these files to the windows machine will vary depending on the type of database you have. If you do not know which MultiValue Basic functions to use, look for OSWRITE, WRITESEQ, or Q pointers that allows you to write to a specific directory location.

If you are on a Linux, AIX, or *nix system, you'll likely need to setup SAMBA or FTP to be able to write to a specific Windows directory.

(See the web site to find examples of writing text and sequential files to the O/S.)

Getting SMTP Up and Running

If you haven't used the SMTP service before, it must be running on your IIS server. The SMTP service is part of the default IIS installation. If you performed a custom IIS installation and chose not to install the service, you can easily add it from Control Panel Add/Remove Programs. To do so, open Add/Remove Programs, click Add/Remove Windows Components, select the Internet Information Services check box, and then click Details.

In a default IIS installation, one SMTP server is installed. You can verify that it's running by opening Internet Services Manager (ISM) and viewing the SMTP server. FTP sites appear first, followed by web sites and SMTP sites. If the SMTP server isn't running, right-click it, then select Start.

Smart Host Configuration

Now, let's look at a few of the advanced configuration options. Most large or enterprise companies place SMTP servers behind firewalls that block any direct outbound SMTP traffic through port 25. If such is the case in your shop, your e-mail messages will

end up in the Badmail folder with cryptic router error-message files. If a smart host is available on your network, you can use it to relay all SMTP messages to the Internet.

In most cases, a smart host is simply another SMTP server with permission to relay outgoing e-mail messages from other internal SMTP servers directly to the Internet. Thus, a smart host should be able to connect simultaneously to both the internal network and the Internet to work as the e-mail gateway. If this setup exists in your company, you can simply add the IP address of the smart host on your network and IIS will send its e-mail messages to that SMTP server for relay to the Internet. To set up a relay to a smart host on an IIS machine, follow these steps:

Open ISM, right-click the SMTP server, then select Properties.

Click the Delivery tab, and then click Advanced to open the Advanced Delivery dialog box, which figure 1 shows.

In the Smart host field, type either the smart host's Fully Qualified Domain Name (FQDN) or its IP address surrounded by square brackets (e.g., [10.1.4.25]). The brackets help speed up the DNS lookup process by specifically telling the SMTP server that the numbers are an IP address. Click OK twice to complete the changes.

In the Advanced Delivery dialog box, notice the Masquerade domain field. The masquerade domain makes an e-mail message look as though it came from a point other than its true origin. You can easily fake return addresses on e-mail messages. For example, if I enter "MyDomain.com" in the Masquerade domain field, mail messages would appear to come from addresses such as nathan@MyDomain.com.

Send E-mail

Now comes the easy part. All you need to do is write your e-mail item to the following directory with a unique file name and the extension of ".eml".
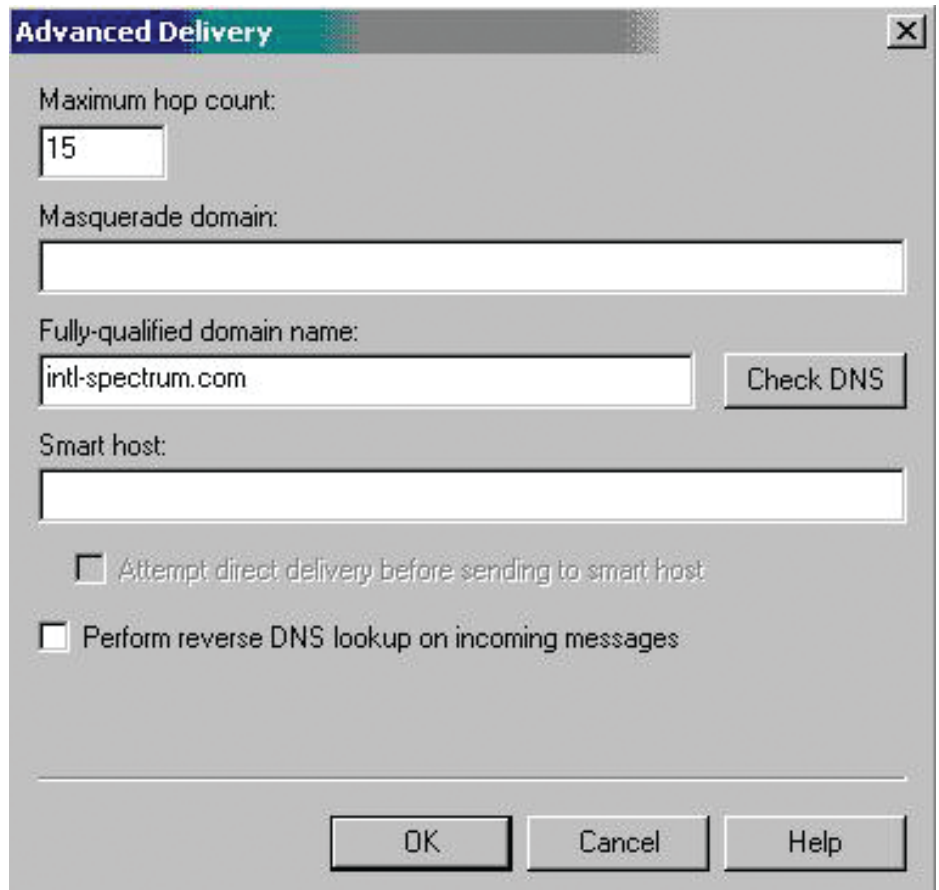
Fig. 1

MultiValue and QuickBooks

Part 2: Exporting to MultiValue

BY MARK JOHNSON

In my first article, I explored how we can export multiple MultiValue data files as transactions into QuickBooks. While QB allows the importing of standard CSV or tabbed files or Excel spreadsheets, it has a private format called IIF whereby a text file can contain rows with multiple table definitions as well as actual data.

In the other direction, it will become evident that QB will appear "short" in the eyes of a seasoned MultiValue programmer. Meaning, while the stock reports and modified reports (memorized) within QB may satisfy a QB-oriented user, your client's requests for more advanced reporting may not be possible within QB.

While there certainly is some form of database within QB, there really is no direct access to it. Their reports drop down and the ability to take stock reports and modify the columns and selection criteria is better than simply stock reports alone. But we MultiValue professionals are used to having full access to our data to manipulate it in any fashion that is necessary to satisfy the requests of our users.

On other systems, you may be able to use Crystal Reports or MS Access to hit

> *We MultiValue professionals are used to having full access to our data to manipulate it in any fashion that is necessary to satisfy the requests of our users.*

the same database that the application is running. You would be careful to not update anything unless you fully understand the application itself. That goes without saying with any "back door" access to the underlying database.

For retrieval purposes only, you don't have to worry about "rocking the boat" with the application. You can cherry-pick whatever fields that you need.

In my decades of slicing and dicing a raw Open Accounts Receivable file within a MultiValue system, I find the QB stock reports adequate and their modifiable parameters to be average. Obviously QB isn't perfect otherwise I wouldn't be writing this article.

There are two separate reports, Summary and Detail, but since we cannot

modify the sort sequence, we are stuck with the sequences pre-programmed into QB. Coming from a full-access database (MultiValue), I get frustrated when trying to create custom reports using this pretty popular file.

In all QB reports there are prompts regarding single or all customers as well as date ranges. You may be able to filter in or out specific transactions. But that's about it, database-wise.

One request in particular was to produce an open AR report by customer and have each invoice within each customer be in date order. The stock Detail report from QB would be in date order only and reviewing each customer from this report is nearly impossible. The Summary report is in customer order but doesn't show any detail.

Virtually all QB reports can be directed to a printer for obvious printing or to a spreadsheet for external review. This is our connection to get the data from QB into regular MultiValue fields for our use.

When you open a stock report, you can see a "modify" button. Stock reports will have a fraction of available output fields pre-checked. You would check

all the fields that you would need. The sort sequence is of no concern with this exercise, although you should insure that the user make the date range as wide as possible to accommodate their request.

Then the user would run the report and send it to an Excel spreadsheet. Now we take over with any popular Multi-Value tool to import a spreadsheet into a standard MultiValue data file.

Once imported, now the world is our oyster. We can sort and break-on with standard MultiValue commands or create highly specialized reports using MultiValue Basic. In either case, the user gets the best of both worlds. The only issue would be the timing between the two systems.

There are inherent reasons that accounting departments use QB or other external systems. My two cases have a few reasons of their own, and it's hard to debate from our perspective.

First — The incoming pool of clerical talent is much younger and certainly would not have experience with a traditional green screen environment. Their experience with anything GUI, especially the Microsoft experience, will naturally cause them to type certain expected commands, <ctrl-S> for Save, <ctrl-P> to Print etc. Even if you have converted many of your long-term MultiValue apps to GUI, there may be the chance that it doesn't follow the Microsoft model.

Second — Along that Microsoft experience, virtually every command and function has a keyboard shortcut. As a novice MS GUI programmer and experienced Accuterm GED programmer, I barely add them as I usually don't have time for them until moving to the next project.

Third — Companies like QuickBooks have invested zillions of dollars and man-hours into perfecting the user interface and application, despite being somewhat limited in complex database reporting. You can't argue with the extent that QuickBooks is installed.

Fourth — For most companies that use MultiValue, the accounting department is not as important as the departments that align with the company's business function. Whether manufacturing, production, inventory, distribution, logistics, or other business functions, most companies that use MultiValue are not accounting companies, ie. CPAs, banks, brokerages etc. Thus, their programming dollars over the years has gone to core business improvements and not helping those in the accounting department. This leaves the door open for QB or other larger networked accounting environments to come in.

Sidebar — From a recent conversion away from a 20+ year highly evolved MultiValue full company system to Great Plains, the accounting portions (AP, AR, GL) were outstanding and made the bean counters happy. The manufacturing and inventory modules of Great Plains were very "plain" (pardon the pun) and when the full conversion was complete, 80% of the company felt it was a step backwards from their highly tuned MultiValue application. C'est la vie.

No one asks us whether to retire part or all of our beloved MultiValue applications in lieu of a shiny new (albeit more plain) system. We are usually informed when it is too late to make a counter offer.

Thus, we should embrace the opportunity to manage both environments, knowing that we still have plenty of work on both systems. **IS**

MARK JOHNSON is an independent MV consultant with over 30 years experience on virtually all platforms. He specializes in maintaining legacy systems, helping his clients maximize their investment while participating with current technologies. When the time is right, he facilitates migrations to the proper contemporary platform. He can be reached at majprogramming@gmail.com or 908-601-4052.

FROM THE PRESS ROOM

Revelation Software Announces the Release of OpenInsight Development Suite 9.2.1 and the Universal Driver 4.7

Revelation Software announces the availability of the latest release of OpenInsight Development Suite (OI) 9.2.1 and the Universal Driver 4.7. OI 9.2.1 provides the following new components:

- OpenInsight for Web (O4W v1.1) – A web 2.0 development toolkit that empowers OpenInsight and MultiValue developers to rapidly create browser based forms, reports, menus, dashboards and programs.
- Source Code Management (SCM) – With this new feature, source code can be grouped into modules and saved as a unique version as the code is compiled.
- Enhanced System Editor ++ – Our system editor has been redesigned to improve speed and functionality.
- U2 DotNet Connector – The U2 Connector has been redesigned to utilize the UniObjects.Net library instead of the InterCall libraries. This redesign allows the implementation of connection pooling.

"The response to our O4W v1.0 product contained with OpenInsight 9.2 was so well received that we placed the majority of our development efforts into the extension of O4W," said Mike Ruane, President/CEO Revelation Software. "Our development community asked for many new features to be included in the next release of O4W, and we delivered!"

The Universal Driver 4.7 and Universal Driver 4.7 (NUL) have been released in conjunction with OpenInsight 9.2.1. The Universal Driver 4.7 replaces the Universal Driver 4.6 and will be offered in two versions. The free bundled version supplied with OpenInsight 9.x network user licenses and the standard version to be used with OpenInsight 8.x and below as well as Advanced Revelation.

"The UD 4.7 (NUL) will now be supplied at no cost to all current members of our Works Subscription program", said Robert Catalano, Director of Sales, Revelation Software. "MultiValue and SQL developers can purchase OpenInsight as a front end development toolkit without our Linear Hash database at a cost of $60.00 per user with a minimum of ten (10) users for browser-based applications. Included are both the O4W web development toolkit and the UD 4.7 (NUL) software."

OpenInsight Development Suite 9.2 is available from Revelation Software, or through their network of resellers. See Revelation.com for details.

About Revelation Software

Founded in 1982, Revelation Software delivers a suite of application development tools and companion services that take full advantage of leading network computing architectures, messaging, groupware, and client server platforms. Today, the company's flagship product OpenInsight is the only database development and application environment that provides both Windows and Java-based GUI s tools to develop and deploy web-based and client server applications that support native and relational XML, SQL, Lotus Notes and the leading legacy Multi-Value data sources such as Arev, Pick and IBM Universe. There are more than 1.5 million licensed users of Revelation products across 60,000 deployed sites worldwide. The company has offices in Westwood, New Jersey, as well as a European distributor in the United Kingdom, and an Asia Pacific subsidiary in Australia. ■

Synergetic Data Systems, Inc. Introduces Click-2-View

Synergetic Data Systems, Inc. (SDSI)has announced the release of Click-2-View, our Desktop Client (DTC) application. Click-2-View is an application that allows the UnForm Document Management Solution (DMS) to be integrated with any application accessed from a user's computer without API integration. Click-2-View can be configured to offer the user the ability to select a single document or search for any group of documents relevant to the information displayed on the application screen. Click-2-View can be setup quickly and can be used with more than one application on a user's desktop.

Version 8 of UnForm is a platform-independent software product that creates, delivers, stores, and retrieves graphically enhanced documents from ERP application printing. A complete, end-to-end Document Management Solution (DMS), UnForm interfaces at the point of printing to produce documents in various formats for printing and electronic delivery. Exact replicas of these documents can be stored in a secure archive for later retrieval by users via a web browser interface, the new Desktop Client, or via a powerful command line API, allowing document access to be built into any application. External documents, such as scanned invoices or signed delivery documents can be imported into archive libraries for further workflow automation.

Version 8 of the UnForm Document Management Solution (DMS) has been widely received and implemented. UnForm has been deployed in over 9000 companies in 15 countries.

SDSI products add features, functionality, and value to hundreds of commercial software products and custom applications. Laser forms, PDF document generation, electronic document delivery, document archiving,report writing, and unique control over Microsoft Office are just a few of

the many capabilities found in our software.

SDSI products are sold through a network of hundreds of resellers, OEM's, and distributors, as well as directly to independent end user customers. SDSI was founded in 1988, with the mission of creating unique and powerful software products. Today our software is used in major industries such as accounting, distribution, retail, manufacturing, healthcare, and in hundreds of other industries.SDSI is the developer of the UnForm Document Management Solution, as well as other productivity tools for UNIX, Linux and Windows.

Contact: Dan   Schmitt Synergetic Data Systems, Inc. 800-446-7374   Ext   125 www.unform.com ■

DesignBais International Announces Support for the TigerLogic D3 v9 .NET API (MVSP)

DesignBais International is proud to announce support for the TigerLogic D3 v9 .NET API (MVSP). This support will allow D3 users of DesignBais to use the native D3 v9 API to connect between the DesignBais Web component and the DesignBais Database component instead of a third party component. This connec-

tion will increase throughput and decrease the points of failure.

For additional information please contact DesignBais International at info@designbais.com or call +1 877-889-9777. ■

Marc Kahn Software LLC Announces the Multi-Value Phantom Job Processor

Marc Kahn Software LLC is pleased to announce the product launch of the Multi-Value Phantom Job Processor (MVPJP).

MVPJP helps to increase the efficiency of your operations by scheduling and automating jobs to be run in the background. It runs as an application within your MVDB environment and is designed to be managed by non-technical personnel. It is cost effective, user friendly, and highly reliable.

The core functionality of MVPJP was developed in 2005 as a tightly integrated add-on to Columbia Ultimate's "The Collector System", meeting with a high level of acceptance in that marketplace. The original software has now been modified to run on any MVDB platform which supports background processing.

In addition to running the scheduled jobs automatically, MVPJP offers:

- A Phantom Console program for schedule maintenance and monitoring.
- An audit trail of historical events.
- Email notification of process completion and other phantom events.
- Flexible query screens to assist in managing the background tasks.

If you would like your computer operations personnel to "work smarter, not harder" then we'd like to talk with you about adding MVPJP to your IT toolbox.

About Marc Kahn Software LLC

Marc Kahn has been active as an independent software contractor since 1982. Early on he was introduced to the Pick database system, which has been the focus of his software career since the beginning. Over the years, he has provided custom programming, consulting, and product development services for a wide variety of clients.

For more information:
call 360-571-2232,
email marc.kahn@mvpjp.com,
visit http://mvpjp.com,
or chat with Marc at the International Spectrum Conference ■.

BlueFinity mv.NET 4.2.2 Enhancements

In accordance with the BlueFinity roadmap, the development team at BlueFinity International has released a new version its

extensive toolset designed for creating Microsoft .NET- based applications requiring access to MultiValue logic and databases, mv.NET 4.2.2.

mv.NET enables software designers to combine the power and flexibility of proven MultiValue technology with the feature rich Microsoft .NET environment. Its seamless integration with Visual Studio provides a potent rapid application development environment fine-tuned for the MultiValue developer.

As BlueFinity continues to roll out groundbreaking enhancements to its flagship product, its developers are closely watching the market's future directions to bring the latest technologies to the MultiValue and .NET developer communities.

"The on-going development work on the mv.NET product range demonstrates BlueFinity's determination, commitment and ability to provide the MultiValue community with industry leading development tools for the Microsoft .NET environment," explains David Cooper, Lead Developer at BlueFinity.

"These tools allow the rapid creation of state-of-the-art, industry standard applications which can be deployed across any of the wide range any of today's interface devices — workstations, browsers, handhelds, phones and other mobile devices."

mv.NET 4.2.2, available now, includes the following enhancements:

- D3 R9.Official support for D3 Release 9 and mvBASE

Release 3 plus support for the new MVSP socket connectivity.

- SSH-based connectivity support. Enhancement to support secure connectivity (SSH) over IP based connections for all databases supported by mv.NET.

- Allow multiple EMR locations to be defined and listed within Entity Modelling in Data Manager. A convenience enhancement to allow multiple EMRs to be accessed within the Data Manager.

- Support for the Silverlight AutoCompleteBox control. An enhancement to the XAML Wizard to allow it to generate code incorporating the standard Silverlight AutoCompleteBox control. This allows dynamic lists (representing data selections from the underlying MultiValue database) to be displayed in a control similar to that used in web browsers to capture the desired web URL address.

Plans for the March release of mv.NET 4.2.3 are well underway with the following advancements included:

- Support for jBASE's jRemote connectivity option. Enhancement to the connectivity layer of mv.NET to allow it take advantage of the new secure sockets connectivity feature within jBASE Release 5. This provides increased connectivity performance for mv.NET applications running against jBASE 5.

- C# code generation from Solution Objects. Full support for Solution Objects to generate C# code as an alternative to Visual Basic.

- XAML Wizard chart designer. New functionality to the XAML Wizard to enable data visualization charts to be easily defined and incorporated into applications.

- Auto-hide of data form controls based on security. Enhancement to the Silverlight run-time support to allow controls on a form to be automatically hidden/disabled based on user security privilege.

- XAML Wizard multiple language translation support. Enhancement to the XAML Wizard to allow foreign language translations to be defined for all literal strings used within Navigation (menus) and data forms. This allows menus and forms to dynamically change their language based on the user's locale setting. Screen content is automatically resized to take into account string length differences across different language translations.

- Allow sharing of entities across models. Enhancement to the Entity modeller in Solution Objects to allow an Entity defined in one model to be referenced and used in another.

- Group entities together into logical collections within Entity Modelling and XAML Wizard. Enhancement to Solution Objects to allow grouping of Entities into logical collections within the

Entity Modeller and the XAML Wizard.

For more information:

Contact sales@bluefinity.com for more information or register for a free trial or online demonstration at www.bluefinity.com. ∎

FusionWare Has New Branding: dbLynx and mvLynx Brands

When we first started as Liberty Software, back in 1992, our flagship product was the ODBC driver for PICK. This was the first ODBC driver released outside of Microsoft, and was clearly targeted to the MultiValue Database community. We rapidly ported this to support almost every version of MultiValue database out there, including, at the time, R83, Universe, Unidata, Alpha Microsystems, Sanyo/ICON, and many more.

As we worked with our customers, we were often assisting them as they integrated their MultiValue data with many other environments, including SQL Server, Oracle, DB2, Sybase, Progress, and many other databases. Often, we had to do the research to figure out how to make these environments integrate well, and even how to optimize the data that we had transferred over. The result of this is that we began to develop expertise and tools to work with many other platforms

and databases.

Today, we have a wide suite of products, tools and services that assist our customers in a wide range of data middleware, BI, portal, mobile and other areas. Some of our products have value to both mainstream and MultiValue customers. This has made it difficult to market these products without creating confusion.

As a result, we are announcing two new brands to allow our customers and partners to better identify the products, services and bundles that serve them best:

mvLynx for MultiValue

mvLynx identifies products, services and bundles that are uniquely targeted to the Multi-Value community.

dbLynx for Mainstream and Relational

dbLynx identifies products, services and bundles that provide value to customers who work with our Microsoft, IBM, Oracle and Open Source solutions.

About FusionWare

FusionWare Integration is a leading solutions provider specializing in business intelligence, integration and middleware products and services.

We have been delivering leading edge integration and business intelligence solutions for more than a quarter century. The name may have changed from Liberty to FusionWare but the same great team continues to innovate and deliver solutions to companies of all sizes.

From our early days as one of

the first providers of ODBC database connectivity solutions under the name LibertyODBC, to providing solutions in web services, e-commerce and business intelligence, our strong belief in standards and providing value to our customers has made us an industry leader.

Our MultiValue data connection products support everything from the original ODBC to JDBC to newer technologies involving Java, .NET and SQL Azure and are available for every flavour and version of PICK. The Integration Server product is based on platform independence and standards based connectivity. Web Services and XML are key foundations to the FusionWare Integration Server.

Around the world, FusionWare products and services have gained a loyal following among Fortune 500 companies and organizations of all sizes, including Elizabeth Arden Salons and MeadWestvaco Calmar.

For more information, visit http://www.fwic.net/ ■

Datatel Appoints John Shea as Chief Marketing Officer

Datatel, Inc. today announced that it has appointed John Shea as chief marketing officer. In this position, Shea will focus on a range of product and strategic initiatives to facilitate the company's next stage of growth. He most recently served as chief marketing officer of Rivermine, a leading provider of automated solutions that enable organizations to control their communications expenses.

Shea brings over 20 years of technology experience to Datatel. He has a consistent track record of rapidly scaling businesses. At Rivermine, he was responsible for overseeing corporate and product strategy, product marketing/management, demand creation, mergers and acquisitions, training, and corporate communications. His efforts resulted in the successful acquisition of the company in late 2010.

Before Rivermine, Shea served as vice president of product marketing, management and corporate strategy for Nuance, the leading provider of speech solutions worldwide. Shea joined Nuance in 1998 during its early stages and was a key member of the management team that helped to grow the company from four live customers to over 1,000 during his six-year tenure. He was heavily involved in Nuance's highly successful initial public offering, which resulted in a market valuation of over $5 billion.

Earlier in his career, Shea worked for Intel, Inc., and was a member of its Pentium Processor marketing team that helped to grow product shipments from zero to over 10 million units per quarter in two years. He also has held marketing and technical positions at OnLive Technologies and Booz, Allen & Hamilton.

Shea holds a bachelor of science degree in engineering from the University of Notre Dame and a master of business administration in marketing with highest honors from Northwestern University's Kellogg Graduate School of Management.

"John is an extraordinary addition to our executive team," said John F. Speer, Datatel president and CEO. "His exceptional track record in successfully growing technology businesses, deep understanding of scalable marketing processes, and keen focus on customer success will help us deliver outstanding solutions to our customers and accelerate our strong financial results."

"I'm excited to be joining the Datatel team," said Shea. "The company has built strong customer relationships and delivers leading-edge products to the higher education marketplace. I look forward to helping colleges and universities realize new levels of value with Datatel's solutions."

About Datatel, Inc.

Datatel is a provider of innovative technology products, services, and insight to higher education. Colleges, universities, and technical schools across North America partner with Datatel to build Strategic Academic Enterprises dedicated to achieving student success. The company has focused on higher education since 1979. For more information, visit www.datatel.com. ■

First Steps to Securing Your Information in the Internet Age

BY BRYAN SHUMSKY, REVELATION SOFTWARE

At the end of 2010, the Gawker Media consortium of blogs was hacked. User accounts and personal information for many (or all?) registered users were taken. This was particularly embarrassing for Gawker, which counts amongst its properties Gizmodo, the popular tech blog, the opinion on the internet was that they "should have known better."

It seemed, a few years ago, that you couldn't open the newspaper (or watch TV news, or view online news sites) without seeing some story about compromised corporate databases in the headlines. Lots of big companies — credit cards, banks, credit reporting agencies — suffered from, and succumbed to, attack.

Other than the Gawker incident, there hasn't been much talk about this in the news lately. Have the attacks stopped? Most industry experts don't think so. Nor do they believe that all IT systems have somehow become so secure they are no longer vulnerable. Instead, companies have just stopped reporting such attacks, in an effort to keep up consumer confidence. After all, would you want to deal with a bank that had just lost your account information?

Our MultiValue installations typically present much smaller targets than, say, Chase or Wells Fargo, but I'm sure that none of us would relish the thought of telling our customers, "Sorry, someone has stolen your personal information from my computer system." Yet odds are we are as vulnerable to these attacks, or more, as the Big Guys. What can be done to make our systems more secure?

Recognize You Have A Problem

Does anyone access your computer, either physically on a console, remotely via a network, or over the Web? If the answer is "yes," then your system is vulnerable. Even if you only use RS232 serial connections (remember them?) to old-fashioned terminals (remember them?), if someone has physical access to your system, they can physically extract the data. Modern PC-based systems will have USB ports, but even older dedicated MultiValue machines had floppy drives. As the number and types of connections to the outside world increase, so does your exposure; if your system is accessible from the Internet, there's the possibility that anyone in the world can reach it.

Minimize Your "Window Of Exposure"

The first step towards securing your system is to minimize the amount of confidential or compromising data that you store, if possible. Many of our applications accept confidential data, such as credit card numbers, and store them away as part of data records without thought as to whether they might be used again. This kind of data should be cleansed from the system.

Of course, there is also confidential data that cannot be removed from the database; a perfect example of this kind of data is the user password. In these instances, our goal must be to keep the data useful to our application, but useless to anyone accessing the data improperly.

Making Data Both Useful And Useless

In order to achieve this goal, our main weapon is encryption — specifically, a type of encryption called "hashing." A hash is a one-way encryption. When data is appropriately hashed, it is impossible (or very difficult, and hopefully beyond the means of the majority of hackers) to decrypt and find the "clear text" original information. This means that you cannot ever "recover" the original data, which is sometimes inconvenient; but it also means that loss of that data does not endanger the end user's account. For the remainder of this article, when discussing encryption in general, we are specifically referring to hashing of data.

While all of the MultiValue databases encrypt the password used to access the fundamental user accounts, many applications have additional passwords that application developers build into their login and authorization processes, and often these passwords are stored in clear text. Even though you may not consider any of the data in your application "confidential," the fact that many users will re-use the same password for many different activities means that you should do your utmost to keep the password secure.

A first pass may be to encrypt the password using the built-in encryption "user exit" or algorithm provided in your MultiValue database. Then, when validating the user during the logon process, rather than comparing the text they have typed in with the clear text password stored in your data files, your application must first encrypt the text that the user has typed with the same algorithm, and then compare the encrypted value with the encrypted password stored in your data files. Your application will have the same level of authorization security as before, but now the data stored in the system is useless to others.

Securing Browser Access

An additional complication arises if the application is available over the Internet via a web browser. It is possible for the data sent from the browser to the host to be intercepted and interpreted, and thus any passwords could be comprised. One solution is to use the "https" protocol, which encrypts *all* the data traveling between the browser and the web server, but this introduces a performance penalty, and if the only "secure" data is the password, it may be overkill.

An alternative (in those instances that don't require full "https" security) is to encrypt the password *in the browser* before sending it across to the web server. In this way, the "plain text" password is never transmitted, and thus cannot be intercepted; the encrypted value is stored in the database, and it is

> *Do not attempt to create your own encryption algorithm, or hashing scheme. Unless you are a trained cryptographic analyst, it is unlikely that any "home grown" algorithm will provide strong security.*

the encrypted value that is again compared before access is granted to the system.

In fact, it is possible to enhance this scheme so that it provides additional protection against anyone trying to "spoof" your system — for example, someone intercepting the encrypted password and then transmitting that to

your server so that they could pretend to be an authorized user.

OpenInsight's web product, O4W, uses this enhanced method to protect the user data and prevent "spoofed" authentication. The basic+ routine that generates the login page builds a script that includes the current date and time, and encrypts those values — along with the user-entered password — before sending the information back to the server (fig. 1 and fig. 2).

Notice also the inclusion of a variable called "Salt". A salt is used when generating an encrypted value to make the value more unique. When storing the passwords in the database, O4W combines the clear-text value and the salt (uniquely generated for each O4W site), so that even the same password will generate different encrypted val-

Fig. 1 Generating the form

Fig. 2 Behind the scenes: Encryption "event" on button

FIRST STEPS TO SECURING YOUR INFORMATION IN THE INTERNET AGE

Continued from page 21

ues on different systems. The salt, in combination with the date and time, are combined with the clear-text user entered password, which are then all sent back to the server in a single hash value.

The host, when receiving the encrypted values, validates that the date and time are within acceptable limits (five minutes, in our example), and then combines the date, time, and previously stored encrypted value, to make its own hash. If the hash passed in from the browser matches the hash generated from the data on the server, the user is validated and authorized to continue (fig. 3 and fig. 4)

Don't Compromise The Security You've Built

O4W uses a standard encryption algorithm (sha1), though different imple-

```
Case CtlEntId _eqc "btnlogin"
    O4WResponse()
    * evaluate what the user entered
    uName = O4WGetValue("o4wNAME")
    pwd = O4WGetValue("o4wPWD")
    bIsBad = 1
    userperms = ""
    If bDoEncrypt Then
        otime = O4WGetValue("o4wTime")
        salt = O4WGetValue("o4wSalt")
        hash = O4WGetValue("o4wPWDHash")
        currdate = DATE()
        currTime = time()
        date.part = Field(otime, ".", 1)
        time.part = Field(otime, ".", 2)
        If date.part = currdate Then
         * ok to proceed - same day
        End Else If date.part = currdate - 1 Then
         * maybe just passed midnight; adjust our time part so that it looks like it's a negative time in the
        time.part = time.part - 24*60*60
        End Else
         * can't be valid
        time.part = ""
        end
        If time.part <> "" And (time.part < currTime And time.part + 300 >= currTime) then
            * within 5 minutes - ok to proceed
            If salt = ourSalt Then
                * salt matches - correct for our system
                userPerms = Function(@auth_proc(uname, hash, otime))
            end
```

Fig. 3 Retrieve and "sanity check" passed in values

```
rslt = ""
If uname <> "" And pwd <> "" then
    * Check against permissions file
    permRecord = Repository("ACCESS", @APPID<1>:"*UINFO*O4W*":uname)
    CALL SET_STATUS(0) ;* OK IF THIS DOESN'T EXIST
    * hash together (already hashed) password from the file with the original form's login time
    HVAL = O4WI_ENCRYPT(permRecord<1>, ptime)
    If hval = pwd Then rslt = permRecord<2>
End
```

Fig. 4 Behind the scenes: Generate hash in authorization routine

mentations may use different algorithms. However, it is critical that the algorithm used be 'strong' — resistant to attempts to crack it. Do not attempt to create your own encryption algorithm, or hashing scheme. Unless you are a trained cryptographic analyst, it is unlikely that any "home grown" algorithm will provide strong security.

If you are planning to implement the security scheme described above, you must make sure that the encryption algorithm is available to use in the browser as well as on the MultiValue platform. The built-in encryption "user exit" or algorithm will probably not be available; instead, you may be able to implement one of the MD5-based algorithms in your MultiValue application, and include a comparable javascript function to perform the encryption in the browser (O4W uses the same javascript library from within Open-Insight as it includes in the browser, guaranteeing that the calculation will be identical).

Once the data has been encrypted on your system, as mentioned above, it is impossible to recover a "lost password." Your application should be designed to allow for an administrative override, if applicable, or the ability to generate a new random (temporary) password for the end user.

While encrypting your data might eliminate the possibility of the convenient, old-fashioned "Email your lost password?" option, it also reduces the risk that undesirable third parties can get access to your users' data — which is the biggest inconvenience of them all. **IS**

User Interface Design

De Facto Standards and Unicorns

BY SUSAN JOSLYN

Q: How many software engineers does it take to change a light-bulb?

A: None. They just change the standard to darkness.

It's easy to change the standards. That's because few companies actually have any kind of documented standards in any area, least of all around the user interface. Paradoxically, this area where you will find the least definition in standards is the highest profile. It's where the user meets the machine, what the software looks like, how it behaves at the most quintessential level. Yet we tend to let it grow wild, like a patch of weeds. A lot of companies are searching for that perfect face lift — some sleek new graphical interface that is going to bring a smoothly polished, consistent, and professional feel to their software. But we expect it to happen organically. Plant the new interface, add water, and we will find ourselves in a neatly ordered world filled with de facto standards and, just as likely, populated with lovely white unicorns!

De facto standards — isn't that an oxymoron? What do we mean when we talk about standards? Sometimes we use the word standard to talk about a bare minimum. An item is "standard" as opposed to "premium" or "upgraded". Sometimes it is an inclusion word — this car comes "standard" with cruise control and tinted windows. Standards are often associated with regulations imposed by some governing body. That will usually happen when there are safety issues or potential for

What the software does and how it does it can increasingly be separated from how it looks. A different look and feel can be dropped over an existing application without affecting the nature of the application itself.

fraud. There isn't much of a safety issue in the user interface of a typical business application. You can't even get a paper cut from a bad user interface! What we are probably looking for in user interface standards is uniformity — "User Interface uniformity." And yet, understandably we are motivated to avoid the sameness with other applications. We want our software to have our creative signature. Plus, the reality is that most companies have a conglomeration of different purchased and home-grown applications along with those that have been acquired via company mergers and acquisitions along the way. If you could document hard and fast rules about how the software should look — even if the target weren't constantly moving — how would you go about getting all of the different programs to comply with it?

We've been getting away with following the leaders and flowing with the current. In reality, that natural flow of currents doesn't wind through green hillsides of frolicking unicorns. Rather, we find ourselves in eddies and brack-

ish backwaters. As lovely as our vision of organically grown graphics may be, it is probably better to have a plan. But like any of the other things that we know we could improve, taking time out to develop and document standards takes precious time away from actually developing the applications and using them in business. This is compounded by the frustration of documenting standards just in time for them to be made obsolete by some flashy new front end.

This is where the latest generation of graphical user interfaces is going to take us to the next level. What the software does and how it does it can increasingly be separated from how it looks (fig. 1). A different look and feel can be dropped over an existing application without affecting the nature of the application itself. A decade or so ago when the term "business rules" sprang into use it was meant to convey the separation between what the software did and how the software
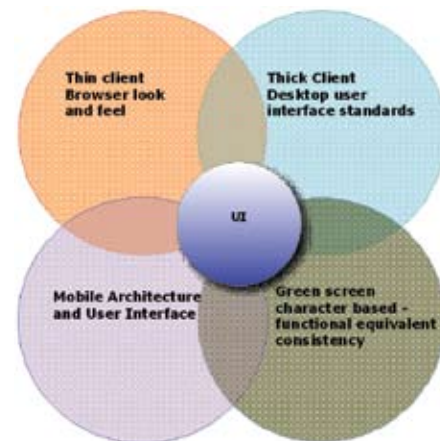

Fig. 1

looked. Now we are moving to the next layer with this new term "lookless controls." We can define not only the business rules, but the interface rules. The "controls." Maybe they are "logic controls." The actual interface can be slid over those like a slipcover. This idea of lookless controls is going to be the game-changer for user interfaces. We can create a corporate skin that we put over purchased and homegrown applications alike. We can even give users choices for their own skins, and — here's the important part — building a design standard at that middle layer will enable us to adapt to upcoming interfaces, like mobile phones and, well, golf bags, vending machines, and coffee makers. We form a layer of standardization that comes between what our application does and how it communicates with the human.

If we harken back we can remember when the global acceptance of open standards for communications between computers revolutionized the way we work and play. Take e-mail for example. In the early days, e-mail was only available between similar servers and applications. So one could only e-mail to and from the same application over a specific communication protocol. As soon as communication standards were published and adopted, suddenly anyone could connect to anyone — regardless of what was sending the message and what was receiving it. E-mail, EDI, XML, SOA (service oriented architecture). The list goes on and on, all dealing with computer-to-computer interfaces. Computer-user interface standards have lagged behind. So it is this latest abstraction layer — this ability to separate the most superficial components of look/style from the user interface — that will enable us to develop meaningful standards for that interaction. Standards that will allow us to move more quickly rather than holding us back.

It is still a lot of work. Even trying to work out what you need to work out is a workout. But you may be convinced that the time has come. If you go out

looking for guidance on UI standards, be prepared to read a lot of opinions and find very little consensus. Still, there is excellent food for thought out there. Blogs where people have taken the time to document their own experiences — successes and failures — can save you some time. There are books that offer guidance from every angle — psychological, artistic, functional, and statistical. The wiki on Eclipse user interface standards offers some excellent high level guidelines. How can you go

wrong following the advice offered in figure 2?

Guideline 1.1

Follow and apply good user interface design principles: user in control, directness, consistency, forgiveness, feedback, aesthetics, and simplicity.

Fig. 2 Guidelines from Eclipsepedia

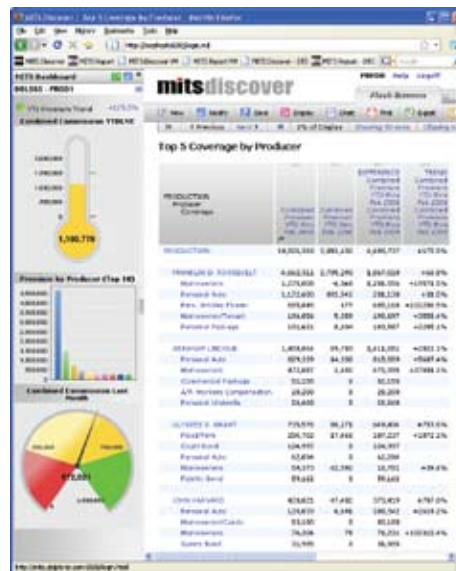The best exercise may be to start with your current application — green

SENDING EMAIL FROM YOUR MULTIVALUE PROGRAMS – PART 2 – USING IIS

```
c:\inetpub\mailroot\pickup
```

As easy as that sounds, there is one catch. If you write directly to the folder, and the e-mail is large, you may find the SMTP server trying to process the e-mail before you are ready. So instead of writing the e-mail directly to that folder, write it to a temp folder. You can then move the e-mail from the \tempmail folder to the \pickup folder. You can do this easily with the following command-line process:

```
move c:\inetpub\mailroot\
pickup\tempmail\testemail.
eml c:\inetpub\mailroot\
pickup\ /Y
```

Again, doing this can be easy or hard depending on if you are on the same machine as the SMTP service. The Windows-based MultiValue systems have it easy in this regard, since you can just use the EXECUTE statement to process a native Windows command.

Take a look at figure 2 for an example program of how I've done this. The program was written in D3/NT, so you will have to look at your MultiValue system's documentation to see how to do this using other systems.

The key things you will need to replace are how the directory file pointer

```
SUBROUTINE SEND.EMAIL(TO.ADDR,FROM.ADDR,SUBJECT,MSG)
EQU AM TO CHAR(254), VM TO CHAR(253), SVM TO CHAR(252)
*
OPEN "EMAIL.PICKUP" TO EMAIL.PICKUP.FILE ELSE
    *** Can not open file, then create the file pointer
    *** in Master Dictionary.
    *
    Q.ITEM = "Q"
    Q.ITEM<3> = "DOS:c:\inetpub\mailroot\pickup\tempmail\"
    OPEN "MD" TO MD.FILE ELSE STOP 201, "MD"
    WRITE Q.ITEM ON MD.FILE, "EMAIL.PICKUP"
    *
    *** Try opening the file again.
    *
    OPEN "EMAIL.PICKUP" TO EMAIL.PICKUP.FILE ELSE STOP 201, "EMAIL.PICKUP"
  END
*
EMAIL.DATE = OCONV(DATE(),"DWA")[1,3]
EMAIL.DATE = EMAIL.DATE :", ": OCONV(DATE(),"DD")
EMAIL.DATE = EMAIL.DATE :" ": OCONV(DATE(),"DMA")[1,3]
EMAIL.DATE = EMAIL.DATE :" ": OCONV(DATE(),"DY")
EMAIL.DATE = EMAIL.DATE :" ": OCONV(TIME(),"MT")
EMAIL.DATE = EMAIL.DATE :" MST"
*
EMAIL.ITEM = ""
EMAIL.ITEM<1> = "To: ": TO.ADDR
EMAIL.ITEM<2> = "From: ": FROM.ADDR
EMAIL.ITEM<3> = "Subject: ": SUBJECT
EMAIL.ITEM<4> = "Date: ": EMAIL.DATE
EMAIL.ITEM<5> = "MIME-Version: 1.0"
EMAIL.ITEM<6> = "Context-Type: text/plain"
EMAIL.ITEM<7> = "Context-Transfer-Encoding: 7bit"
EMAIL.ITEM<8> = ""
EMAIL.ITEM<9> = MSG
*
EMAIL.ID  = DATE():TIME():".eml"
WRITE EMAIL.ITEM ON EMAIL.PICKUP.FILE, EMAIL.ID
*
MOVE.CMD "move c:\inetpub\mailroot\pickup\tempmail\": EMAIL.ID :" "
MOVE.CMD = MOVE.CMD :" c:\inetpub\mailroot\pickup\ /Y "
EXECUTE \!\: MOVE.CMD
RETURN
END
```

Fig. 2

USER INTERFACE DESIGN DE FACTO STANDARDS AND UNICORNS

screen or whatever it is — and think about developing standards with *two* interfaces in mind. In addition to the general guidelines mentioned on Eclipsepedia, here are a few things to keep in mind as you get started.

- Make the standards simple, practical and implementable.

- Make it easy for a developer to find the particular standard or guidance that he or she needs.

- Make each standard/definition an entity that can be mapped across multiple UI's.

- Make sure it is intuitive. If you aren't sure, borrow someone's seven year old and see if they can figure it out.

As in all insurmountable tasks, it will help to break it down. Four main components to start with are: Navigation, Forms, Reports, and Documentation. They cross-apply to one another, too, but you have to start somewhere. Error messages and help messages fall under documentation in this broad categorization. *Some people* might argue that they deserve their own category and should be addressed first!

Navigation — How do you get into the application and then move around within it? Switchboards, drop down menus, command buttons on forms? Do you make one choice at a time or are many options presented at once? How does that translate to the application running on the green screen? How about on your phone or mobile device?

Forms — How are the forms laid out? How is navigation handled within the form? Borders, fonts, and colors might be something that you later allow to be changed as a skin. In fact, it may be a good idea to decide in advance which aspects you do plan to be skin-changeable. Think about the "hardware."

How are forms shaped? Aligned? Do you use interior boxes? White space? On small devices you may want to get away from people *entering* an answer and rather to let them choose one of several choices. How will those choices be presented?

Reports — First decide what qualifies as a report. Are reports even necessary anymore? Is a dashboard or graphical representation going to be the output of choice instead? To keep from junking everything up with a lot of different things, some early decisions should be made about when each form of "informational output" would be best used. Guidelines about context or the type of information and in what format it should be presented. Then get into the standards about how each will look, including traditional reports.

Documentation — Let's face it, in the new age of user interface, people are not going to spend much time with their face in a book. We want the answer to the question we are facing *right now*. We want it quickly, and we don't want to have to move very far. So documentation almost disappears. In its place well designed help and error messages will be the wave of the future. To be successful, these must convey a lot of information in a small amount of space with few words. Establishing some standards for how that will be accomplished is arguably the most important aspect of designing a modern user interface.

Oh — and don't forget to add the unicorns. **IS**

SUSAN JOSLYN is the President of SJ+ Systems Associates, Inc. and is the author of PRC?, a complete, integrated software development life-cycle management / IT Governance tool for U2. She has worked with U2 (nee Pick/Multivalue) and SB+ software the beginning (both hers and its) and has specialized in IT Governance, including quality, compliance and life-cycle productivity issues since the early 1990's.

CLIF NOTES: THEY'RE EVERY-WHERE! (WELL, NOT QUITE YET)

some by choice. So paper is still a fact of life. Still, the advantages of digital cannot be ignored.

So as of the last issue, you will note that each article in the magazine has a QR code that links you to the online version of that article. See an article that you want to read on the train during the commute home? Scan the QR code and bookmark it. Leave the magazine at the office. Frankly, we are experimenting, so watch for other uses as we test the waters.

Some of our advertisers have started including QR codes in their ads. They can provide you with more detailed product information or demonstration videos. It's kind of a hybrid., sort of a paper hyperlink.

We are also using them at this year's conference. Since I am writing this in advance of the conference, I don't want to go into any details and spoil the surprise. But for those of you who aren't able to join us, we will let you know what you missed in our Recap article in the next issue.

So what are QR codes used for? That is primarily limited by your imagination. Sometimes at the various evening functions at the conference, such as the Sneak Peek evening, some form of entertainment is available, like a magician specializing in close-up sleight-of-hand. I offered one of my own suggestions for this year's conference. Unfortunately, the management of both International Spectrum and the hotel were in agreement that an event having a no-host bar, QR codes, and a tattoo artist was a bad combination. **IS**

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**

SENDING EMAIL FROM YOUR MULTIVALUE PROGRAMS – PART 2 – USING IIS

works, and how to execute a Windows command.

Drawbacks to IIS

The main drawback to using IIS is that you have to have access to the Windows machine to be able to write its directories. This may create a security concern in your enterprise.

The other issue with using IIS is that once it's running, anything that gets dropped into that folder will be sent as an e-mail, and you won't have any control over it. There is no logon process to verify you can send the e-mail through your servers, so if a malware program or user find out that this server exists, then you have a very large security concern.

The next article I will talk about how to send an e-mail using Linux SEND-MAIL. **IS**

BUSINESS TECH: USER OWNERSHIP OF DATA

reasonableness of the data. We can't do the best job without them. Likewise, we can't abandon the data to them and expect success.

Jack Sprat Could Eat No Flack

Another reason for user ownership of data is, frankly, blame. When IT owns the data, we are at fault when it is wrong. We don't generally have the authority to change it, but we have the responsibility and, therefore, the blame. When users own their own data, we might share responsibility or they may own it in total.

In one of my past lives, I was a co-manager for Wenco, a Wendy's franchisee. Part of the management philosophy there was "You don't give the R without the A." In other words, Responsibility and Authority belong together.

It is also useful to remember that if Accounting owns the accounting data and Sales owns the sales data, getting it right is to their credit. If IT helps them get it right, while still respecting that they are in charge, that brings us credit. **IS**

Got a topic you'd like to see in Business Tech? Send it to "Results@Key-Ally.com"

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.

They're Everywhere!
(Well, Not Quite Yet)

I see QR codes.

But unlike the ghosts of dead people, these things are present in our external reality. And I'm not alone. Lots of people see them, including the attendees at the 2011 International Spectrum conference. You see them too. Look up at the top of this page. See that square checkerboard thing? That's a QR code.

The QR "Quick Response" code is a bar code that allows you to store text messages, web site URLs, contact vCards, and other things in a form that can be easily scanned and acted on by a QR code reader. Such scanning software is readily available for smartphones.

Developed by the Denso-Wave subsidiary of the Toyota Corporation in 1994, they have become very popular in Japan. Of course, anything having to do with a mobile phone is always a big hit in Japan, including little electric wheels and leashes so you can walk them next to you like a pet. (Okay, I'm exaggerating. The pets are simply apps that you have to care for or they die.) They are just starting to take hold in other countries, including the United States. If you don't know what they are, you tend to look right over them without even noticing. But as their use increases, you will begin to see them popping up all over the place. So, you might ask, what good are they?

Not a darn thing.

That is, not in and of themselves. But people are beginning to find some interesting uses for them. For example, Google sends out window decals to businesses identified as a Favorite Place. On the decal was a QR code that would link to the business's page where you can read or write reviews, be informed of special promotions available to them, or mark the business as a Favorite.

You can include your contact information in vCard format on your business card. Scan the code and your contact information is automatically added to the reader's address book. Or use it to store a web site address that takes them to your product catalog or a product demo on YouTube.

You may be planning on buying a house. One day while you're out driving around, and you spot the absolutely perfect place with a For Sale sign in front of it. Underneath the For Sale sign there is a box meant for information brochures about the property, a few interior pictures, and who to contact. Of course, the box is empty. Take a close look at the For Sale sign. Some real estate agents are now printing a QR code on it. Scan it with your phone, and the next thing you know, you are looking at all the listing info for the dream house you are standing in front of.

Calvin Klein once put a large QR code on a billboard. When folks would scan it with their phone, it would launch an advertisement that was, shall we say, "spicy." (I bet some of the implications of this kind of use has some attorneys salivating.) But it gets better.

In some regions, McDonald's uses QR codes on its product packaging so people can link to nutritional information about the product they are ingesting. I find it ironic, however, that you have to buy the product in order to get the packaging that has the QR code that gives you the information so that you can determine whether or not you want to buy the product.

And in Japan, some people are putting QR codes on tombs. (What did I tell you?) Scan it with your phone and get information about the deceased.

There you have it. QR codes really can help you see dead people.

Until recently, the discussion of QR codes and smart phones have been closely tied together. Enter the tablet.

One of the problems with digital publishing is that you either needed a notebook/netbook computer or a dedicated e-book reader to take advantage of it in anything other than prose form. I mean, you might read a novel on your phone using your Kindle app (I do), but if we publish a 50 line code example, are you really going to try to read that with that tiny screen on your phone? Not likely. PDF format works for some people, and they even prefer it. That is one of the reasons why International Spectrum magazine is available for download in PDF format. But again, until recently, you needed a notebook or desktop computer to read it.

But the advent of the tablet changes that. Of course, tablets are not as wide spread as phones. And a lot of readers don't have smart phones, some because of the cost of digital plans,

Virtual Conference 2011

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.

Unable to join us in person?

Need to talk with a vendor again?

Missed a session at Spectrum Live?

Spectrum 2011 Virtual:

- Live Chat
- Downloads
- On-Demand Recordings
- Social Media
- Networking

Register Now!
April 12 – 14, 2011

INTERNATIONAL
Spectrum

intl-spectrum.com/virtual/2011