

\$7.00 U.S.

INTERNATIONAL Spectrum

THE MULTIVALUE  TECHNOLOGY MAGAZINE | MARCH/APRIL 2016

A Specification Specification

Also in this Issue:

- Using MultiValue to Manage Multimedia
- Data Replication
- 50 Shades of Interface
- Building a Modern Application: Part 2



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.

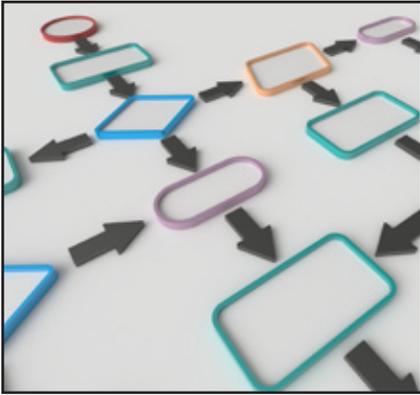
Professional Memberships provide you access to knowledge, solutions, information, and code that you won't find in other locations.

Professional Membership Includes:

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



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



6 A Specification Specification We've all been exposed to project specs. Some barely cover the back of a napkin. Others could crush a small country. It's time to rethink our approach. Bennett talks about specifications, iterative collaboration, and the value of building a better umbrella. **BY BENNETT BAROUCH**

10 Calls May Be Recorded for Quality Assurance: Using MultiValue to Manage Multimedia Too often we think of databases as being just about the data stored in them. Sometimes the database is an index to things outside of it. Multimedia and MultiValue may not seem like a perfect pairing to most people, but when you have to develop a system to sift through thousands and thousands of "calls recorded for quality assurance," a little database architecture can go a long way. **BY CHARLES BAROUCH**

13 Data Replication Unauthorized access to your data is a serious problem. However, blocking authorized people from seeing it is nearly as bad. They can't see it when the system is down or the data is corrupt. Backups are one solution. Data Replication is another. **BY MARK FULLER, NORTHGATE INFORMATION SOLUTIONS**

18 Business Tech: 50 Shades of Interface There's something standing between you and your data. And there should be. Classically, interfaces are thought of as access points, but a good interface is also a speed bump. We need to rethink connectivity. **BY CHARLES BAROUCH**

21 Building a Modern Line-Of-Business Application: Part 2 There have been many new Line-Of-Business applications developed over the past several years. Most of these business software systems either address a missing piece of the overall or are just extensions of the existing monolithic applications that have been around for fifteen to twenty years. This is Part 2 of our article series which explores what it takes to build a Line-Of-Business application from scratch using the tools and features found in modern software technologies. **BY NATHAN RECTOR**

DEPARTMENTS

From the Inside page 4

From the Press Room page 16

International Spectrum and MultiValue are registered trademarks of International Spectrum, Inc. All other registered trademarks are the property of the respective trademark holders.

From the Inside



NATHAN RECTOR
President
International Spectrum

Over the last ten years, new interface technologies have emerged, but LOB (Line of Business) and ERP (Enterprise Resource Planning) applications are not taking advantage of them. This is causing LOB and ERP to become very cumbersome and blocky. I'm not just talking about MultiValue applications; this is across the industry.

The Rules Have Changed for Windows Applications

There is new pressure to update your approach. Windows 10 has arrived on the Business Desktop. Some of us have installed it. Some of us are still waiting a bit longer. Either way, IT departments are working through the user issues of Windows 10 while developers are still wondering how best to use its features. And it does provide new options that we that *should* be interfacing with our LOB and ERP applications. They are not as widely talked about as Universal Apps, the return of the Start Menu, the write-once-run-everywhere design, and the HaloLens Apps, but can become far more important.

If your existing LOB and ERP applications do not fit with the new interface models, then your employees and your customers will start looking for alternatives. People expect the ability to access and update your masses of enterprise-level data from the programs they use the most. They don't want a model which requires them to move from program to program, or device to device, to do their work. That's the old way.

This idea is reflected in Microsoft's new API convergence. It was introduced as part of Windows 10, but really isn't Windows 10 specific. That's the point. This isn't the API

for Windows Desktop OR Windows Tables OR Windows Phones OR Office OR Office 365 OR OneDrive OR Sharepoint OR Azure OR Surface OR HaloLens OR Surface Hub OR iPhones OR Androids OR the slew of other hardware and software technologies on the market today. This is the API for ALL of them. Together. At once.

What this really means is that Microsoft is turning their software and O/S into a set of APIs. Developers will be able to integrate an application into everything that Microsoft supports, instead of building an application that *runs* on top of a Microsoft O/S. These new APIs are based on current standards which span across devices and programming languages. Now developers are even less locked into .NET or C++ or some other Microsoft Language when interfacing with Microsoft tech or software. Windows 10 Desktop is not an O/S anymore. It is transforming into just another UI Framework that interacts and responds to these APIs.

Leverage the Advantages

LOB and ERP applications all have important information-gathering and distribution components. That's what makes them so powerful. Once we integrate these components into this new Windows 10 UI framework, suddenly our LOB or ERP application has more power than it ever had before.

Your applications should all be taking advantage of Windows Live Tiles found in the Start Menu. They should be updating and interfacing with Windows Search to find contact or customer information. They should be interfacing with Outlook directly, and sending and retrieving contact changes from within the Windows People and Outlook

INTERNATIONAL Spectrum

MARCH/APRIL 2016

NATHAN RECTOR
President

CHARLES BAROUCH
Editor

SYDNEY BAROUCH
Editor

TRACEY RECTOR
Layout



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

<http://www.intl-spectrum.com>

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

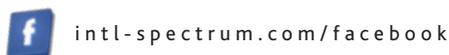
International Spectrum is published six (6) times per year at the subscription price of \$40.00 U.S. in the U.S.A.; \$45.00 U.S. in Canada and Mexico; \$50.00 U.S. for other countries. Single copy rates are \$7.00 U.S. in the U.S.A. and Canada, and \$9.00 U.S. in all other countries. *International Spectrum* is published by International Spectrum, Inc., 3691 E. 102nd Ct., Thornton, CO 80229; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com. Copyright 2016 International Spectrum, Inc. All rights reserved. Reproduction in whole or in part, without written permission, is prohibited.

PRINTED IN USA

NEWS RELEASES/UNSOLICITED ARTICLES

International Spectrum is eager to print your submissions of up-to-the-minute news and feature stories complementary to the MultiValue marketplace. Black and white or color photographs and diagrams are welcome. Although there is no guarantee a submitted article will be published, every article will be considered. Please send your press releases, articles, and queries to: editor@intl-spectrum.com. *International Spectrum* retains all reprint rights.

International Spectrum is a registered trademark and MultiValue is a trademark of International Spectrum, Inc. All other registered trademarks and trademarks are the property of the respective trademark holders.



10 Years Strong & Growing

Integrity, Efficiency, Service,
Guarantee, Global Productivity,
Midwest Rates, Universe,
UniData, D3, Barcode Services,
Nationwide Support, Web Services,
Custom Barcode Programming,
.Net Programming, Mentoring,
Project Outsourcing



www.pickprogram.com
contact@pickprogram.com
(614) 921-9840

Do you... have projects sitting on "the list" and not being completed?

Do you... just need some assistance from time to time?

Do you... want up-front, guaranteed estimates?

Senior Consultants are Currently Available.

"PICK Programmer's Shop has been outstanding to work with. They work with the highest integrity, are very responsive and quite knowledgeable. I can't imagine operating without them."

Patti Rowlette, Rowlette Executive Search

It's not just an anniversary. It's a guarantee.

Contact managers. Notification systems, customer updates, and order completion should all be notifying users through Windows Notification Center. Users should be able to interface with these notifications to interact directly with the needed application or data.

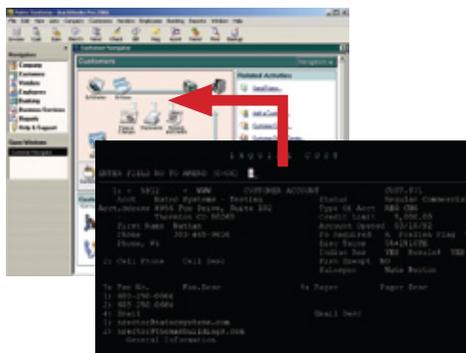
Any data or action presented in those places should also be presented in Mobile Applications on Android and iDevices. The information should be securely accessible from outside the enterprise through Cloud APIs or Cloud Apps.

The API approach brings more power and flexibility to your application, allowing interaction with hardware and software that is closer to seamless. It also permits better expansion and enhancement to the important software: *Your software.*

Stop looking at Windows 10 as just an OS that runs your application. Start looking at Windows 10 as a UI framework for presenting your users with modern features and state-of-the-art integration. IS

mv
QB

QuickBooks API for the MultiValue Database



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

NATEC
Systems



Providing Solutions to your MultiValue Questions

Phone: 303.465.9616

E-mail: mvqb@natecsystems.com

Website: www.natecsystems.com

A Specification Specification

BY BENNETT BAROUCH

Raise your hand if you love spec writing. No one? Raise your hand if you like reading a spec written by someone else. Definitely no one. If you think we don't need to improve this situation, you can stop reading now. If you envision a better world, however...

The first question is the same first question for most things – what are you trying to accomplish? End user documentation? Technical documentation? Release notes? Tutorial?

The first corollary question is the same for most things – who is your audience? C-Suite stakeholders? Consumers? QA? Sales team? Architecture level conference paper? Hands-on information for implementation developers? Professional services customization engineers?

Wow! Are we going to cover all of that in one article? Raise your hand if you would really want to read an article that covers all of that in one pass. Again, no one. Yeah, me neither.

So far we have established a broad topic, and a way to not cover it *because*

We forget at our peril that the development of novel intellectual property is not the same thing as manufacturing. And as an industry, we forget this all over again almost every day. This is a major reason for unreliability in software, and for unreliability in software development schedules.

there would be no audience for that form of expression. This is already better than much of the documentation we have all read.

The Challenge

For this article we will focus on the mother of all specs in the software development world. By whatever name, this spec captures the essence of what the 'product-specifier people' want the 'product-maker people' to produce.

Why not call them Product Management and Engineering?

Because lots of people specify what is to be produced. It's not just product management, but also architects, the CEO, and lots of people who design, implement, and test code, plus people who design data structures and storage, and so on. They are all 'product-specifier people.'

Part of the historical problem with specs is that we want them to do more than we routinely acknowledge. Another part of the problem is that we want them to do more than is actually possible for a written document. For convenience, if not precision, let's call this a product specification.

As you may be aware, the old-school view is that a product specification should catalog every single thing to be done, in unambiguous detail, before any work is begun. That never worked well. You would put a huge amount of effort into creating a massive document that, as the saying goes, "does not survive first contact," either with the development team or with the customer.

That approach misses the hard-won lesson that we do not actually know as

much as we think we know when we begin a project. We actually need specs to take different shapes throughout the development process. To define it all up front is to make a huge investment in documenting something we are virtually guaranteed we will not build. While few companies write specs like this anymore, almost all are still unconsciously committed to this way of thinking: *I'll tell you what to build and you build it.*

This has a horrible track record in the software industry. We forget at our peril that the development of novel intellectual property is not the same thing as manufacturing. And as an industry, we forget this all over again almost every day. This is a major reason for unreliability in software, and for unreliability in software development schedules.

Surely you know that all the cool kids are saying that the only way to proceed

sensibly is to define a little, implement a little, play with it a little, and then do another increment. I think this is both abundantly true and is also dangerously incomplete guidance. If everyone on the team is functioning at the master architect level, and they intuitively share a single, coherent vision of the end result, this might work. Since those conditions do not exist on this planet, we need to broaden our thinking.

In my observation, the claim that good architecture will spontaneously emerge is only true when a team is extremely well-aligned with a far-better-than-average architect (or architecture team). In most cases, a general architecture must be laid out to provide a conceptual framework and direction, and to give life to certain key principles. Otherwise, the result will be dominated by inconsistent APIs, poor performance, low maintainability, over-dependence

on particular engineers, security weaknesses, poor testability, and schedule delays. It will also present increasing difficulty with quality and timeliness on each subsequent release.

The reason we sometimes get around this with an exceptional architect and exceptionally well-aligned team is not because it emerges, but because they actually do these things a little at a time, embedding them in every small and large design decision. Few teams are so good, and gel so well, that you should make your business depend upon architecture spontaneously emerging. Do not over-specify an architecture since it is likely to change during development, but don't expect a pony to magically show up for Christmas either.

By direct analogy, you also need a general test plan, documentation plan, marketing plan, training development plan, and so on. Whatever the prod-



**taking multivalue ...
where it has never been before**

Ladybridge Systems Ltd

17b Coldstream Lane, Hardingstone, Northampton, NN4 6DB, England

Worldwide distributor: Zumasy, 9245 Reasearch Drive, Irvine CA 92618, USA

www.zumasy.com

- **Close compatibility with most other multivalue environments**
- **Easy migration process**
- **Maintenance-free file system for ease of use**
- **High quality documentation**
- **QMClient API for development of GUI and web applications**
- **Low licensing cost**
- **AccuTerm bundled at no additional cost**
- **Many unique features**

www.openqm.com

uct-makers are going to make calls for a *lightweight* plan and lots of ongoing communication. That brings us to a point I want to make exceptionally clear.

What is the Role of a Spec?

A spec begins life as a thought-stimulator and thought-aligner, and ends as a set of rearward-looking revisions explaining what actually came to pass. Most of the success in great team alignment during product development comes from *talking*.

Why can't the spec be the coordination and control mechanism we want it to be? It is not only that we change our minds about exactly what we want as we go, and not only that we change our minds about what we can accomplish in a given time frame. A spec is a low-bandwidth form of communication.

You take all of that inspiration and energy and tone of voice and innuendo and gesturing and eye contact and audience reading and interaction that went on in those great exchanges – that high bandwidth communication that still ran far short of a complete transmission of an identical shared vision – and you reduce it to writing. So, it's kind of, well, stupid to think that a document is going to ensure a no-surprises, no-disappointments outcome.

- Product Specification Point 1: Create a lightweight architecture plan and, as development shows a need to modify it — um — modify it. Do the same thing for test planning, training planning, documentation planning, marketing planning, and anything else that must be done.
- Product Specification Point 2: Expect written documentation to be a memory jogger for things more

fully transmitted in live conversations involving live human beings. Do not expect specs to be magical incantations that cause things to happen in the real world just by being invoked.

Now that we are mindful to not overspecify, we are ready to talk about the bulk of our product specification: feature definitions.

Story Time

Once upon a time, a Customer said to a Sales Person that they wanted an umbrella that was easier to use. The Sales Person told a Product Manager, who went to work surveying the features of umbrellas already on the market. She proposed a new model with a competitively better self-opening handle. The Product Manager told an Engineer to make this super-umbrella, and the Engineer made it even better by making it the most water-repellent umbrella ever. It even had a heating element to make the water evaporate off of it quickly once you came in from the rain.

The Sales Person rushed back to the Customer with great excitement and showed off the new umbrella. The Customer seemed completely uninterested and asked why the Sales Person was showing him this thing. "But you told me you wanted an umbrella that was easier to use – I mean, look at this thing! It's awesome!" "Oh," said the Customer, a bit embarrassed. "I did. But I didn't mean for the rain. I wanted something easier to use for protection from the sun. I found a great new sunblock cream. It's non-greasy and protects from the full spectrum. Once I put it on, I can forget about it and I'm good all day. That's *easy to use!*"

The Moral

The number one mistake in almost every spec ever written is that it specifies a solution instead of a problem to be solved. Not only can you entirely miss the customer's objective, but you limit the search for a solution to a single mind when you could have had everyone in the production chain engaged in brainstorming about how the target problem could be solved.

Be sure to restate to the Customer what you think they have identified as the problem, not the solution they may have proposed. It *is* a good idea to ask them exactly how the solution they propose solves the problem. Make sure you convey all of this to the larger team.

- Product Specification Point 3: Always specify the problem to be solved and start the conversation about how to solve it. Don't end the conversation in the spec's first draft. Make including a specific solution a revision to the spec once the team has figured it out together.

What form should feature specification take? All the cool kids say user stories. Again, this is absolutely correct, but incomplete. Defining a feature is like defining a solution — not so good. Much better to define a utilization of the product. Then we have a richer, more dynamic idea of what the user needs to do, and everyone can add more value. Not, "It stops the car." More like, "As a driver, when I apply variable pressure to the brake pedal with my foot, the car's momentum reduces proportionately, so I can slow or stop in a controlled manner."

The cool kids are right — that is cool. So what's missing? Lots of things. It just does not make sense to express every miscellaneous requirement in this

form, even if you include “non-functional requirements” with each user story (which you should).

We have requirements that apply broadly, and we do not want to repeat them on each individual user story — that would not only be a waste of time, but it would encourage people to not read the specifications we are giving them. Perhaps you have over-arching rules such as: “A response time of more than two seconds is universally unacceptable — where unavoidable, the UI must come back with a message that the task has been scheduled, and notification will arrive when it is completed” or “Security requires all sensitive info to be encrypted with AES.”

This is not architectural stuff, so it doesn’t belong in our architecture plan. It is however very important stuff that belongs in every bit of function-

ality. You should document this kind of thing in prose and ensure that the whole team understands the significance of these factors. Ensure that people in a test role are testing for compliance with these factors at exactly the same priority as “features.”

- Product Specification Point 4: Create a traditional, prose specification for requirements that span individual functional bits (and probably multiple projects). Test for these factors alongside feature-oriented testing. Define specific bits of functionality with user stories. If the “user” is another bit of code, it’s still a call for a user story. The UI is just an API instead.

An additional benefit of breaking it up into user stories is that each bit of work is small enough that people can genuinely understand what is expected

of them, and can learn to give reliable estimates over time.

To conclude, we have many bad spec habits to break and a few simple new habits in which to become proficient. We need to stop thinking of a spec as the final word on anything and stop assuming it is a good form of communication. It’s really just note taking.

When we start making specifications that address problem statements in lightweight, flexible terms and use them to engage the entire team in active discussion, we net a shared responsibility for definition and delivery of a great product on a schedule that is based on reality. **IS**

BENNETT BAROUCH has over 30 years of industry experience spanning design automation for integrated circuits deployed in satellites, financial portfolio software, high transaction volume and big data systems, information management, secure on-site and mobile networking, and IT operations software. His customers run from individual retail consumers to the largest companies in the world.

Groundbreaking work under Bennett’s leadership produced a virtual assistant that could understand 20 million English phrases and respond with a wide array of information and with complete computer-telephony integration. This work was made part of the permanent collection of the Smithsonian Institution, for Outstanding Achievement in Information Technology, 14 years before Apple released Siri and became the basis of the OnStar virtual assistant found in GM automobiles. Bennett has been certified in ITIL, and as a Scrum Master and as a Scrum Product Manager.



can
MultiValue can't do that.

Change your thinking.

Modern MultiValue, LLC
Modern Solutions for Modern Applications
info@ModernMultiValue.com
www.ModernMultiValue.com
+1.225.341.1778

Calls May Be Recorded for Quality Assurance

Using MultiValue to Manage Multimedia

BY CHARLES BAROUCH

The Challenge

Sunrise Credit was facing a technical challenge. They had everything they needed to succeed — the data, the audio files, the technology, the people — but none of it was organized for ease of use. It was my task to take MVON

Express and use it to bring order to the various parts.

The goal was simple: produce a two page website for a select number of internal experts. The first page would allow them to login and trigger a query. The second page would display the results. Sounds easy, right? The devil, as you know, is in the details.

The data for this project was in CSV (Comma-separated value) format. There are tools for treating CSV as a database. There are tools for populating virtually any data-

base from a CSV source. Half a day to load and test... we should be ready for our next task by lunchtime. Except that the CSV source is thousands and thousands of separate CSVs. And, by the way, each one is embedded in its own ZIP file. Oh, and there will be at least a dozen new CSVs every business day.

Of course, since this is a website, we are looking for a fairly fast response. The database needs to be fast to update and fast to query. Over time, the queries will need to get more sophisticated. Looking down the road, reporting, not just interactive use, may be in the cards. We need a built-to-last approach. And we need to develop quickly.

Call Review

Username

Password

Year/Month/Day 2014 / October / 13

Skill ATT

Agent

submit

Fig. 1

Call Review

2014-10-13 [REDACTED]

145 record(s).

New Call Review Request

Play	Download	Account No	Time	Phone	Session ID	Call Result	Agent Result
<input type="button" value="Play"/> 0:00 1:21	Download	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]
<input type="button" value="Play"/> 0:00 0:25	Download	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]
<input type="button" value="Play"/> 0:00 0:43	Download	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]	[REDACTED]

Fig. 2

Feedback

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

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

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

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

PICK PROFESSIONAL

Don't make the mistake of placing your career in the hands of just ANY search firm. We are professionals who understand the importance of CONFIDENTIALITY and RESPECT of a job search, and our database of clients is one of the largest in the country. Unlike the rest, we will work in YOUR best interests' to help you further your career. Because of our dedication and professionalism, we are recognized as the leaders in the PICK/UniVerse/Unidata placement industry in the Tri-State area and throughout the U.S. So if you are tired of putting yourself at the mercy of the rest.

CALL THE BEST! Contact...

Matt Hart

EXECU-SYS, LTD

1411 Broadway, Suite 1220
New York, NY 10018

(212) 967-0505

(800) 423-1964 x 302

Fax: (212)947-8593

Email: mh@eslny.com

Consultants Needed Nationwide

As luck would have it, Dawn Wolthus called me as I was starting the project and asked if I'd like to be part of the first tour of the MVON Express product. The timing was too good to ignore. Here are the details on how we built a mature system for selectively monitoring calls for quality assurance in just a few days.

The Approach

John Ocampo, the head of IT at Sunrise, gave me permission to use whatever tech made sense. This is a rare and wonderful gift. He and Anthony Vidas found me a server, rebuilt it with a reasonably current version of Windows and let me loose.

We installed Apache and PHP using XAMPP. We added the current build of MVON Express. We added MS SQL Server Express. We installed multiple web browsers for testing. We had the tools.

The server had access to two network shares: "Archive" and "Available." Archive would contain the entire library of recordings in pristine form. Available would be the copies I could unzip, rezip, rip apart, extract CSVs from, and organize how I needed them to make the site work at a reasonable speed.

The Unexpected Ally

Perry Stauffer came into the picture at this point. Because we were the very first site for MVON Express, he took a direct interest in watching our progress. It didn't take long for him to trade watching for working. Once his sleeves were rolled up, things moved forward rapidly.

What We Did — All the Tech We Used on the Back End

Windows BAT files are a very efficient way to get a lot of simple work done. We built a BAT unzip, extract

the CSV, and rezip each ZIP on the Available share. That moved us from a flat layout (every ZIP in the same directory) to a functional layout of calltype>yyyymm. Now, to find a call from February for call type ABC, you'd look in S:\ABC\201602\. Reducing the size of each directory by splitting them in a way which matched the query logic (call type is the most important break) made Windows happier — fewer files per folder — and made the storage match the retrieval.

I then wrote a small Delphi program — it could have been batch, but I liked the extra speed of an executable — which grabbed all the CSVs from each directory. Since the CSVs were now outside of the ZIPs, it was a matter of copying a text file to a new directory. This is where MVON enters the picture. The destination directory C:\Ongroup\MVON\RAWINDEX, is visible to MVON as a file called RAWINDEX in the SOUNDWAVE. OS account. We wrote a quick mvBASIC program to collect the CSVs and add them to a single MV database file called CALL_REVIEW. Because this is MVON, we had the option of making CALL_REVIEW an MS SQL Express table. That was the fastest option in terms of building the file and retrieving it.

Perry provided an example of how to make all of that work (See sidebar at <http://www.intl-spectrum.com/resource/685/default.aspx>).

For those keeping score, we now had a single place to go when we wanted to query the contents of any of the CSVs. Because of how we did it, we had the option of pulling the data back using MV or SQL technologies. And both of those options would be quick. That left us with the work of building the

IT audits have you jumping through hoops?



PRC can help you meet your compliance requirements and make IT more agile and productive. No extra work, nothing to remember, nothing to fall through the cracks. Our software development lifecycle tool automatically prevents or detects change according to your criteria. You can deploy, rollback, test and report quickly, automatically and with confidence. Let PRC protect your company's valuable U2 data and software assets.

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



web pages. There are a lot of ways to connect web pages to MV data. There are even more ways to do it with SQL. We chose none of them.

Kevin King taught me a trick during one of his International Spectrum classes. It allows the web developer who is ignorant of MV to still be very productive. We used Kevin's two-bucket solution for this project. Buck-

et one is a Windows directory called REQUEST which MVON Express sees as an MV file called REQUEST. Unsurprisingly, you put your requests from the web in that one. The other is called RESPONSE — in Windows and in MVON Express — and gets the responses back from the database. While I can't give you the code we used for this example, I can give you

```
execute 'SELECT REQUEST'
TODO.rec = ''
done = 0
loop
  readnext REQUEST.id else done = 1
until done do
  read REQUEST.rec from REQUEST.file, REQUEST.id then
    TODO.rec<-1> = REQUEST.id : @VM : REQUEST.rec<1>
  end
repeat

*
loop
until TODO = '' do
  RESPONSE.id = TODO<1,1>
  CRITERIA = TODO<1,2>
  TODO.rec = delete(TODO.rec,1)
  gosub respond
end

*
STOP;* Logical end of program
```

Fig. 3

```
$nugget .= "|" . $YYYYMMDD . "|" . $skill . "|" . $secondary . "|" . $secvalue . "|";
file_put_contents($TheRequestName,$nugget);
sleep(1);
$guardian = 120; // 120 attempts
while (!file_exists($TheResponseName) && $guardian > 0) {sleep(1); $guardian -= 1;}
$TheCSVData = explode('^',file_get_contents($TheResponseName));
```

Fig. 4

code which is fairly close. The example below is fully functional. We just stripped out the project-specific bits [Figure 3].

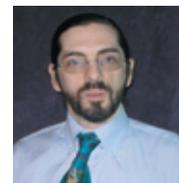
What We Did — All the Tech We Used on the Front End

The site was done in HTML, CSS, Javascript, and PHP. Since the PHP bits handled all of the interaction with the data, let's cover that in a bit more depth [Figure 4].

As you can see, the code isn't very complicated and it is unaware of the underlying database technology.

Final Thoughts

To accomplish this project, we used MVON Express databases, an MS SQL Express database, HTML, CSS, Javascript, PHP, Delphi, batch files, and mvBASIC. That's not counting the O/S, multiple web browsers, 7-Zip, and other assumed tech. We did it with one developer (plus some of Perry's time) and some active support in setting things up. Tech is less and less about using one skill set. Adding MVON Express to this project improved the speed of completion and added to the ease of maintenance. **IS**



CHARLES BAROUCH is the CTO of HDWP, Inc. and the Publisher at HDWPbooks. You can read his writing in

International Spectrum, Theme-Thology, Novo Pulp, PehelionSF, and the Interrogative series, which begins with Tiago and the Masterless.

Data Replication

BY MARK FULLER, NORTHGATE INFORMATION SOLUTIONS

If I were to ask the question: “How do you define what is meant by data security?” I’m willing to place a bet that everyone reading this would initially think of protecting data from unauthorized access, or a more commonly used word, a ‘hack.’

While that is a correct answer, data security is so much more than that. Protecting against ‘hacks’ is extremely important, and in some cases legally required, but maintaining access to data for authorized persons is equally important. It is another form of a data security that must be considered.

Hopefully, we all take backups of our databases that can be restored in the event of a failure. Unfortunately, this can be a slow way to recover a system. And it can lead to frustration from users who cannot perform their tasks without having access to your data. We need more than just backups.

So Just What is Data Replication?

Well, “it’s the frequent electronic copying of data from a database on one computer or server to a database on another.” Furthermore, if it’s *truly* replication, it’s done behind the scenes without you having to worry about it, in as near real-time as possible.

In this article I’m going to explore three options that are available in Reality today that can assist you with data replication: Failsafe, RealityDR, and

Fast Backup. Before we start on these, I have to share a word of caution in respect to relying on third-party snapshot technologies such as those that are part of VM infrastructure deployment for replication. Replication at this level can result in an incompletely replicated database as a Pick system can take several writes to update an item (e.g. out-of-group item) and the snapshot could, in theory, get taken between these writes, rendering the replicated database corrupted.

Fast Backup

That brings me to our first option, Fast Backup. While technically not a replication service, you can use this tool to take a valid snapshot and restore it quickly onto a second database. Reality’s Fast Backup capability allows gigabytes of data to be backed up in just a few minutes and restored to another database just as quickly; far faster than traditional MultiValue saves. The key part of this is that it does have the ability to take a clean, known-state backup that will be guaranteed to be intact. It also boasts the ability to suspend updates to a database so that you can use third-party tools to replicate the database safely, eliminating the multi-write issue highlighted above. You can achieve this with the command `realdump -fs` to freeze updates (and commit to disk) and `realdump -u` to once again allow updates. This command line utility can be embedded

into third-party backup scripts to ensure that a safe snapshot or checkpoint is taken.

A close look at the Failsafe and RealityDR products show some similarities; they are true transaction-based systems and replicate the data in as near real-time as we can possibly get it, depending on the sites’ networking. A true transaction-based system is one that records items that have changed in a transaction log file. It only records the actual items that have changed and not entire file groups or disk sectors where the change has occurred. This potentially saves huge amounts of data-writes as well as affording the users transaction commit, abort, and rollback capabilities.

Failsafe

Failsafe is integrated into Reality and consists of a primary (the live server) and a secondary (the backup server). This feature is available across all of Reality’s supported platforms. Users login to the primary server and perform their usual daily tasks. All updates to files are correctly flagged for transaction logging and are replicated to the secondary server. The update operation is also recorded in a Rawlog partition or file. This holds all before-and-after images of the items.

Once the update(s) are committed, the ‘after’ image is transferred to a transaction clean log file (commonly called a

CLOG file) while simultaneously being transferred to the secondary server to ensure that the update is replicated. Certain operations do not result in an item being transferred but instead result in an operational record being logged, such as a DELETE or CLEAR-FILE operation. Index updates are not transferred because the indexes will automatically be updated on the secondary server upon replay of the transferred item onto that database. [Figure 1] explains this further.

Configuration and management of Failsafe are performed via the utility TLMENU. In the event of a primary failure or planned maintenance of the primary system, the administrator would switch the networking for users to the secondary server. They'd use TLMENU to swap the flags, making the secondary into the primary. The administrator can tend to the now isolated old primary server, make the

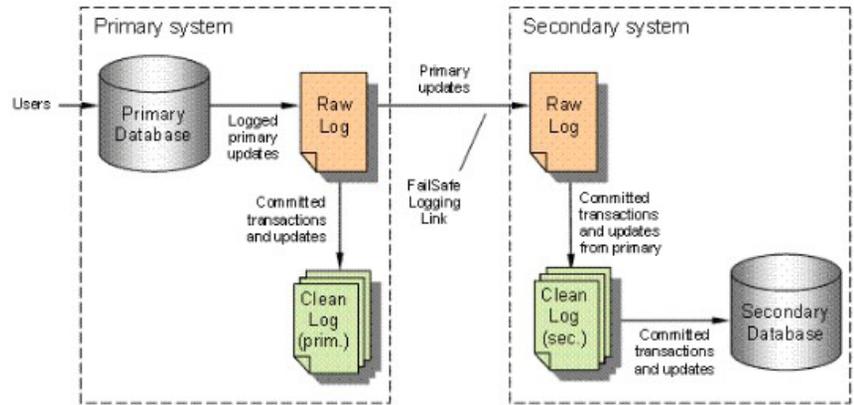


Fig. 1

necessary repairs or maintenance, and then once again use TLMENU to re-synchronize the servers. All the while, the users remain logged-in, performing their daily tasks.

As the servers are typically identical, but not necessarily so, the administrator may choose to leave the users on the new primary. This is all done without loss of transaction integrity and with

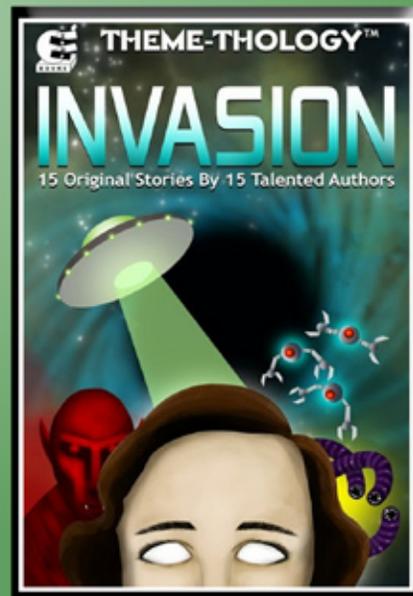
minimum loss of service availability to users.

Depending on the throughput of the database, tens of millions of items can be updated each and every day. This, in turn, can generate significant network traffic between the primary and secondary server. For this reason, it's recommended that the two systems have an isolated network between them which is used purely for the fail-

THEME-THOLOGY: INVASION

Voices by Lisa A. Kramer
 I Was a Teenage Alien by LJ Cohen
 Singularity by Jeremiah Lewis
 Not Like Us by Mike Reeves-McMillan
 That Kind by Charles Barouch
 Yellow by Bill Ries-Knight
 An Invasion of Ideas by Jeremy Lichtman
 Famine, with Fries by Jefferson Smith
 The Several Monsters of Sainte-Sara-la-Noire by Michael Williams
 Going Viral by Rachel Desilets
 Dead Planet Scrolls by Timothy Hurley
 Red Vapor by Michaela Susanne
 The Worms Crawl In by Michelle Mogil
 Nano Nation by CM Stewart
 The Woods, The Cellar, and Cover art by Aaron Wood
 All other interior art by Juan Ochoa

WWW.THEME-THOLOGY.COM



safe traffic. This would lead to the best performance, with updates being replicated in real-time on the secondary server.

Failsafe *can* operate over the general user population's network. This should be configured as a backup route for traffic should the dedicated route fail. Due to the network traffic that this will create, it's recommended that these servers be in close physical proximity to each other.

While I'm on the subject of performance, it's worth mentioning that you should carefully consider the placement of the Rawlog and transaction logs (CLOG's) on a server as you would wish to try to avoid conflicts at the disk level between these and the actual database(s). If at all possible, I recommend having the database, Rawlog, and CLOG files all on separate disk spindles to gain the best performance at the hardware level. Again, this is dependent on the application, since lighter applications might not need the extra speed advantage.

The utility `tlmenu` is extremely comprehensive and impossible to document in this article. I would highly recommend that you refer to the user guides to assist in navigating through the options.

RealityDR

Reality Disaster Recovery provides a further level of replication capability for Reality databases, including off-site replication of a standalone system [Figure 2] or an off-site replicated system for a fail-safe system [Figure 3]. This further strengthens the security of your data.

Increasing the resilience of a fail-safe system gives the added benefits that if the secondary is taken offline, you

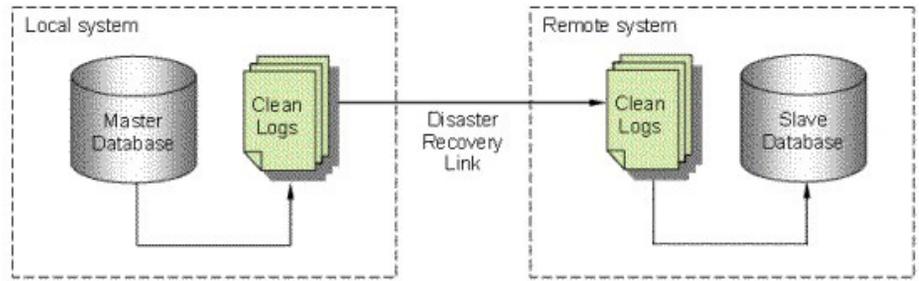


Fig. 2

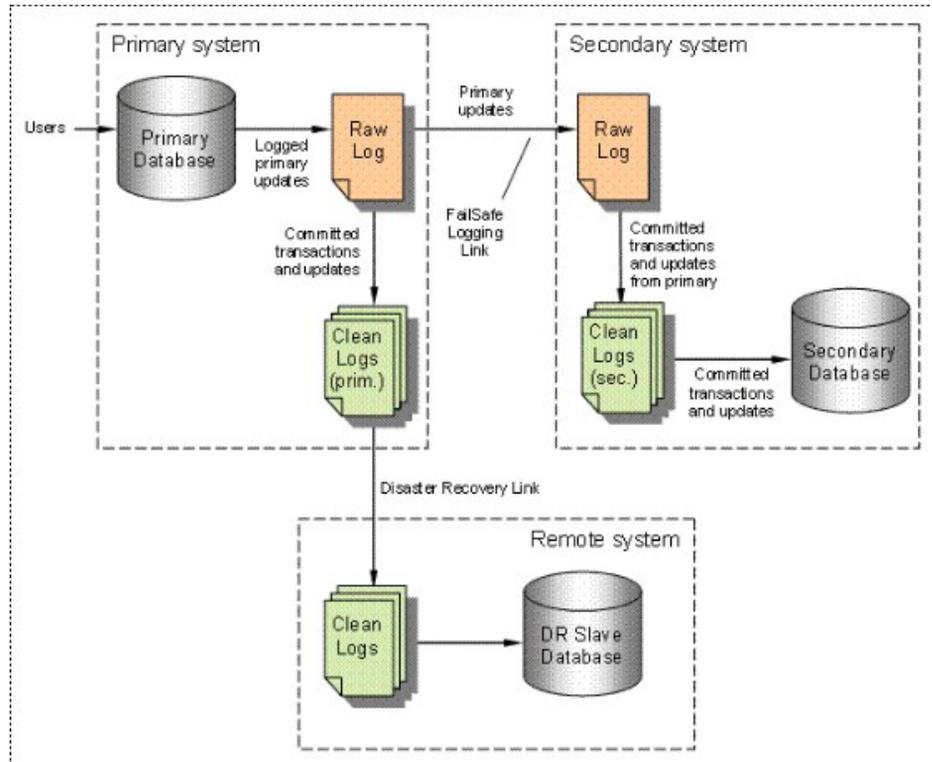


Fig. 3

would still remain protected against a primary failure. Furthermore, in a true disaster situation where both primary and secondary systems are unavailable, an off-site backup system ensures that your data is secure.

The above diagram [Figure 2] shows a typical stand-alone RealityDR system while the diagram below it [Figure 3] details RealityDR together with Fail-safe.

RealityDR is designed to operate over a loosely coupled network, i.e. the server is physically located some distance away such as in a Disaster Recover Center and/or where the network link may not be as reliable or less perfor-

mant. The main server takes the role of the master with the slave system (RealityDR) kept up to date by means of the CLOG files copied from the master. If the slave system becomes unavailable for some reason, the master (the one users are connected to) continues as before; transfer of the clean logs will resume when the slave again becomes available. All of the operations are managed via the utility `tlmenu`. Activation of the slave system, should this become necessary, is a manual step.

It is good to have multiple options for your data security because business needs are always changing. **IS**

FROM THE PRESS ROOM



Entrinsik Informer Named a 2016 Top Rated Business Intelligence Tool by Software Users on TrustRadius

Entrinsik Inc. has recently announced it has been named a 2016 "Top Rated" business intelligence tool by software users on TrustRadius (www.trustradius.com).

TrustRadius gives software buyers access to objective, in-depth user reviews and is the leading site for business users to share their experiences with various software providers through detailed, categorized reviews, lending valuable insight to prospective buyers with real-world use cases. The Top Rated Badge is recognition of best products in a category based on user reviews and ratings. The badge is awarded to those products that have an above average satisfaction rating within the Business Intelligence Tools category.

"Entrinsik Informer is a Top Rated product on TrustRadius," said Megan Headley, Research Director at TrustRadius. "Users particularly like

that it enables organization-wide data ownership, has strong reporting capabilities, and is easy for end-users to learn."

Since the first review on March 25th, 2014, Entrinsik Informer has received over 70 user reviews and ratings evaluating the products' capabilities and ease of use.

"The Informer reporting tool is easy to use for users of all levels of technical expertise, and users can create reports without help from IT. Security administration is very flexible, and dashboards allow immediate access to statistics and key performance indicators."

- Kara Bosch, Director College Admin Support Training at Central Piedmont Community College

"I find Informer a great product that will compliment any public safety agency by giving them the ability to create their own tactical reports with tools that are very user friendly. End users are now more educated regarding software applications. It just seems like sense to put the right tools in front of them to enhance their work environment." - Joseph Riquelme, Senior Technical Solutions, Consultant at Motorola Solutions

See all reviews of Entrinsik Informer on TrustRadius at www.trustradius.com/products/entrinsik.

Entrinsik Informer is recognized as being an innovative solution in data management

providing access to real-time data from multiple sources, multidimensional analytic capabilities, and intuitive data visualizations, all accessible anytime, anywhere from an end-user friendly web-based interface. ■



Ladybridge Releases 3.4-5 of QM

This release introduces the ability to log every command executed on the system together with its origin. This may be useful for diagnostic purposes but has been added primarily to meet audit rules applicable to some applications.

The DIR.SEL.OMIT.HIDDEN mode of the OPTION command causes select operations against directory files to omit hidden items.

The ED, MODIFY and SED editors now all support automatic sequential record id generation.

The SCAN command provides a simpler and more logical semantics for the DISPLAY mode of the SEARCH command.

The AS clause of the SETPTR command has been extended

to allow use of an alternative hold file name.

The SHOW.LIST command displays an active numbered select list, allowing the user to remove items from the list using keystrokes based on those of the SHOW command.

The SP.VIEW command has been extended to allow hold file items to be deleted.

Two new QMBasic functions, BOOL() and NUMERIC() have been added as an efficient way to transform data to Boolean or numeric form. The STR() function with the second argument set to one (or omitted) provides the most efficient way to force a value to be stored as a string.

The QMBasic FCONTROL() function has a new mode, FC\$CRLF, to allow a program to override the default operating system dependent selection of directory file line terminator.

The QMBasic JBUILD() function has been modified to encode any mark characters in a JSON string element as hex Unicode values. The effect of this is that a JSON string can always be stored as an element of a dynamic array such as a database record.

The QMBasic MATBUILD and MATWRITE statements have been enhanced to return the last used matrix element index via the INMAT() function in a manner similar to MATPARSE and MATREAD.

FROM THE PRESS ROOM

The QMBasic SET.ECS.MAP.NAME() function can be used to change the ECS character map within an application program.

The sort system now allows data or keys over 32k characters.

The QMClient API has been extended to allow use in multi-threaded applications on all platforms (previously Windows only).

This release introduces the Python wrapper for the QMClient API. See the qmclient.py item in the SYSCOM file for details.

A new tool catalogued as *FIXDIR has been added to provide a simple way to correct the line terminator in a directory file item that has been moved between Windows and Linux in binary mode. ■



Rocket Releases wIntegrate Version 6.4.0

Rocket has just released version 6.4.0 of Rocket wIntegrate, a user interface for MultiValue databases, which adds enhanced features for IoT, file sharing, and security. The new release also allows for improved user experience with comprehensive integration between character applications and the desktop.

A key feature of version 6.4.0 of wIntegrate is support for IPv6, which provides the ability to create more IP addresses for years to come.

In addition, with this release, users can now share files within *and outside* of their organization with wIntegrate 6.4.0's integration with Rocket R/Link for MultiValue. R/Link is Rocket's secure file sharing solution that offers key security advantages by residing on the network or private cloud.

And finally, with the new release, users are no longer susceptible to the known exploits for Secure Socket Layer (SSL) and Transport Layer Security (TLS) with the addition of the new OpenSSL 1.0.1mj library which is included when users upgrade to wIntegrate 6.4.0.

wIntegrate 6.4.0 is available for the following operating systems: Windows 7, 8, 8.1, Windows 2008 R2, and Windows 2012 OS, and is certified for use with Rocket UniVerse, Rocket UniData, Rocket D3 DBMS, and Rocket mvBase DBM. ■



New Features Announced for jBASE 5.5

jBASE 5.5 will be the first major production release since Zumasy purchased jBASE in

2015. Freshly released in early April, version 5.5 will include the following new features:

System Manager

jBASE System Manager provides graphical access to important system administration data, which is historically trapped at the command line of your Pick system. Now you can easily configure and monitor backup jobs with hashfile verification; enjoy built-in visualization of system performance; enable alerting around critical system processes; and monitor license usage. With System Manager, you can even add licenses on the fly, without having to do a system restart. Keep track of key data, critical processes, and license usage with jBASE System Manager.

Integrated jRCS

jBASE 5.5 also includes our powerful Remove Communication Services at no additional cost. With JRCS, new hires can choose their favorite development tools and languages and directly access your MultiValue apps and data. It's a practical approach that preserves and leverages your proven apps and data while allowing new development in virtually any language or framework.

Audit Logging

Protecting the data in your MultiValue system has never been more important and now you can do it for free with jBASE 5.5. An integrated feature of the new release of jBASE, Audit Logging allows

you to create an audit trail of activity on chosen application files. It's easy to setup and administer. And it adds very little overhead to the running of your system. Best of all, it's free and it gives you the flexibility and feature-richness so that it can be customized to meet your unique requirements.

Case Insensitivity

Migrating to jBASE just became easier. jBASE now supports case-insensitive applications. You can run files, queries, and commands by simply turning on the jBASE case-insensitive switch.

Spool2PDF

jBASE 5.5 includes a framework for spooling print jobs to nearly anything. Popular targets include PDF, email, and web services but virtually any conceivable target is possible. The framework allows for your user-written code to be called during the spooling process to filter, extend, and transform spooled output and to deliver it to nearly any target. With this new feature, it is now simple to add exciting new capabilities to your applications without modifying the underlying application code.

Simple Installers

You can install and configure jBASE 5.5 in just a couple of clicks. New migration tools are also available to help you easily import your existing MultiValue system. Click, install, and you're done!

50 SHADES OF INTERFACE

BY CHARLES BAROUCH

I'm working with a lot of — compared to me — young people lately. Several of them are in school for IT and related careers. Talking to them has helped me to see certain things with new eyes. Recently, I ended up in a conversation about APIs which quickly became a conversation about interfaces in general. This article is an expansion of that conversation.

In Yer Face?

The question sounds naive: "What's an interface?" Proposed definition: *An interface is a method for allowing a person or process to interact with our processes and data.*

That answer is really broad. It has to be. APIs, for example, are interfaces. That puts massive Java VM API — see the Oracle vs. Google lawsuit — in the same category as `connect_db` in PHP. APIs aren't the only interfaces. So are, for example, web pages which have a database behind them. It's not out-of-scope to call the Amazon.com website an interface to their order system. We can make the argument that a web browser's address bar is an interface to the entire internet.

Saying X is an interface is analogous to saying X is music. A toddler twanging a single string might make something

*Impenetrable is undesirable...
insubstantial is also
undesirable. Interfaces,
therefore, live in the vast
middle ground between those
two widely spaced polar
opposites.*

we would have to admit is music. But the term is so broad as to also include a marching band, a symphony orchestra, and Ozzy Osbourne.

Notepad: Your New Database Solution

Instead of infinitely expanding the idea, it might be more productive to come at the answer sideways. How about a test case? Try this one: Can we call Notepad a complete database solution? Short answer: Yes. The long answer? Keep reading.

Can I create a new file using Notepad? Yes! Replace the entire file? Yes! Update it? Yes! Delete? Yes!

That's CRUD, the core functions of any database management system. What about query? Let's see: Open a file in Notepad and press Control-F. What about multi-table query? I can open several files at once and I can

visually scan from one to another, so *kindalsorta*. Reporting? Cut-and-paste. ETL? Cut-and-paste again.

So, yes, in the most literal sense, Notepad *is* a valid database interface. To understand why it isn't, we need one of two things: (A) an ounce of common sense, or (B) a better definition.

Amended definition: *An interface is a method for controlling how a person or process interacts with our processes and data. There is a reasonable assumption that we don't want them to have unrestricted access.*

The Stuff Lining the Walls of the LHC

Impenetrable is undesirable. The Notepad example reminds us that *insubstantial* is also undesirable. Interfaces, therefore, live in the vast middle ground between those two widely spaced polar opposites. We don't want them opaque and we don't want them transparent.

Here are some other words you may need to add to the definition:

- **Industry-Standard** – Some businesses have traditions and long-standing rules for what an interface should allow or deny.

- **Legal** – Banks, Women’s Shelters, Police Departments... these are just a few examples of organizations who are under legal obligations to control the degree and kind of access permitted.
- **Audited** – Wikipedia, for example, provides an editing interface. As a Wikipedia editor, I can change nearly anything. However, my actions are monitored. People review the changes, and who made them, for appropriateness.
- **Redacted** – There are interfaces that have automated censors. You can get the documents you want, but only the version which has been scrubbed of identifying information.
- **Graduated** – Different roles and security clearance levels are baked into the mechanisms. This is how you can write one interface while still giving Corporal Smith less information than Major Jones.
- **Time-Sensitive** – The Freedom of Information Act is just one example of “secret today, public tomorrow.”

Times Change

A million years ago, when I started working in computers, green screens roamed the Earth, and the screech of the three-hundred baud modem presaged dawn of the era of work-from-home. Back then, I could make a strong argument that a non-networked, non-graphic, user-unfriendly system was a complete database solution. So what if it stored passwords unencrypted? No way to reverse-lookup phone numbers? No worries. *Sixty-megabyte* disk drive limit... why would you need more? Things have certainly changed.

On a scale of Notepad to whizz-bang, where does your system sit? Can we trade data in multiple formats via a properly secured access? What about in-app, context-sensitive help? Are your transactions running within commit/rollback envelopes? And, can I see you dashboards?

TCL is an interface. That doesn’t make it an enterprise-ready solution for delivering reporting capabilities. Assembler is a programming language but that doesn’t make it as easy as mvBASIC or Python.

I used to chat with a fellow at ADDS named Jim Roberts. He told me that the first accounting software he created was written with a soldering iron. Times change. Are you changing with them?

Shiny

That doesn’t mean that there is no place in the world for TCL, or Notepad, or soldering irons. It doesn’t mean that your upteen-year-old app is automatically bad. And let me go on record as saying that a system with a million interfaces is not ideal. There is, one hopes, a balancing point between making tally marks in soft clay and rebuilding your infrastructure every six to eight months.

Also: MultiValue is not the opposite of shiny and new. Our database vendors are constantly innovating. There are people who go to work every day, building modern Multivalued software. There are new 4GLs, reporting tools, compliance software, ERP systems, and the list goes on and on.

So?

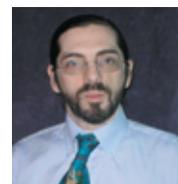
This started with a college kid asking me about tech. He’d like to know more about MultiValue. It isn’t old to him; it’s just another way to get things done.

Why does he see it as viable? Because it works seamlessly with the automation tools he’s using. Because data flows from it to big-screen TVs on the walls around him. Because the people who pay his salary run a business — a responsive, robust business — using it.

MultiValue is only old and dull if we work hard at making it appear that way. Nathan’s series on building modern applications is an example of new ideas with a MultiValue core. Have a look at the Press Room in this issue, or any issue. Come to a Spectrum conference, or a Zumapalooza, or the next Rev conference, or maybe a MultiValue University. Look at the websites of the people who support our industry.

Take a little time one afternoon and download the last two years worth of this magazine and *read*. Then Google the authors of the articles. You may be surprised just how much time and effort is expended — by consultants and system houses — to find new ways to do... well... everything.

Some of us may be old but our software, and our way of thinking, don’t have to be. If you need me, I’ll be updating the new BeBackBy app, which has an OpenQM backend. Or maybe I’ll be working on the MVON-powered Soundwave call-monitoring software. Wherever I am, whatever I’m doing, I’ll be learning something new. **IS**



CHARLES BAROUCH is the CTO of HDWP, Inc. and the Publisher at HDWPbooks. You can read his writing in

International Spectrum, Theme-Thology, Novo Pulp, PerhelionSF, and the Interrogative series, which begins with Tiago and the Masterless.

MARKETPLACE

ACCOUNTING

Natec Systems

www.natecsystems.com | nrector@natecsystems.com



QuickBooks API for the MultiValue Database

- Read/Write Directly to Quickbooks Databases
- mvQB API is Designed for the MultiValue Program to Use
- No Need to Learn the Internals of QuickBooks
- QuickBooks Pro/Premier/Enterprise



NATEC
Systems

Providing Solutions to your MultiValue Questions

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

COMPLIANCE

SJ+ Systems Associates

www.sjplus.com | sjoslyn@sjplus.com

CONSULTING

Drexel Management Service

www.drexelmgmt.com | dconboy@drexelmgmt.com

Execu-Sys, LTD

www.eslly.com | mh@eslly.com

HDWP

www.HDWP.com | results@HDWP.com

Modern MultiValue, LLC

www.ModernMultiValue.com | info@ModernMultiValue.com

PICK Programmers Shop

www.pickprogram.com | brian@pickprogram.com

Precision Solutions

www.precisionline.com | Kevin@PrecisOnline.com

DATABASE

Ladybridge Systems Ltd

www.ladybridge.com | sales@Ladybridge.com

REPORTING

Brian Leach Consulting, LTD

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

TERMINAL EMULATOR

Zumasys

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



AccuTerm® software, the leader in terminal emulation, allows you to access your MultiValue application—whether it is on-premises or in the cloud—from any Windows device.

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



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

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

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

Modern Line-Of-Business Application

Part 2

BY NATHAN RECTOR

In last issue's overview, we discussed the lack of major changes in LOBs (Line-Of-Business) applications despite the wealth of changes in consumer application design, hardware interfaces, and user expectations. In this article we'll talk a bit more about what goes into creating a modern Line-Of-Business application from scratch.

Logging

The problem with logging is that it is almost always added as an afterthought, not as part of the core structure of the application. It shouldn't be. Logging is essential. It doesn't matter if all you are building is a mobile app that displays pretty pictures, a consumer X-Box game, or a Line-Of-Business application. There are times when you must have additional information to help find issues when the users start calling about a problem.

In our area of focus, Line-Of-Business applications, logging is even more important because of the complexity of the application and data. Line-Of-

User and Applications Logs: This is where the debugging information belongs. The when, where, and why belong here.

Business applications require different types of logs for different purposes:

1. System Errors/Information
2. User and Application
3. Security
4. Change

Each of these types requires different pieces of information to be stored. Likewise, they need different notification options during updates. For anyone that's written log files before, you know that the process generates massive amounts of *rarely-used* data. When files are needed, it is very important to have all the clues, or at least the details needed to recreate the original data if something is corrupted.

System Error/Information Logs

This type of log is often over-used as a catch-all. The piling-on of anything that isn't specific to other logs can hide the important parts under mountains of miscellany. Part of the problem is that people try to use System Error / Information Logs for debugging. They are better when used for tracking specific events that are happening in the application. i.e.: *Backup Started, Resize Done, Low Memory, Low Disk, Background Process Started, Workflow Engine Failed.*

System Logs should have the following structure to handle the right type of information:

ID – Unique Record ID

<1> – Date

<2> – Time

<3> – Port or PID

<4> – Account/Database ID

<5> – User ID

<6> – Process/Program creating the message. This is used to help track all

messages being generated by a specific process. For example: a backup or file-resize process.

<7,n> – Error Code(s). Error Codes that represent standard, human-readable, short messages. This is similar to the use of *201* or *404* errors used with the MultiValue ABORT command.

The error code should contain more information than just an error number. For example, if this log contained the “201 - file not found” error message, then the error code may look something like this:

201~CUSTOMER

Or if you were recording missing record IDs with the “202 - item not found” error message, then we'd include the file name (CUSTOMER), the ID that was missing (1234), and the associated file/record (ORDER and 554-1) that was connected to this missing data:

202~CUSTOMER~1234~ORDER~554-1

<8> – Human Readable Short Message. This is a short summary that is human readable to explain the error. Try to keep it to one line of information

<9,n> – Human Readable Long Message. This is a detailed summary of the error, and/or how to go about fixing the error. There may not always be information here, but this is so a human knows where to go and what to do if this error shows up and something needs to be done.

Example: DAT Tape Read Only. Eject the DAT tape from drive, and check to make sure the Read-Only switch is not set. Reinsert the tape, and go to Menu options 1, 2, then 5 to restart the backup.

<10,n> – Stack/Trace information. If this is a major error, then there should be tracing information. This information is all about how the process got to the point where we needed to log the error. It may be the subroutine CALL stack, or it may be a programmer's description of the steps that got to this point.

User and Applications Logs

This is where the debugging information belongs. The *when*, *where*, and *why* belong here. User and Applications logs need to be able to be turned on and off at will, and should be considered transitory, not persistent. Since these are transitory logs, they do not need to be saved in backups, and should be cleared on a regular basis just due to the amount of space they use. Most application developers create these logs in their client applications but neglect the server side user and application logs. We need them in both places.

User and applications logs should have the following structure to handle this type of information:

ID - Sequential ID :”*”: User ID :”*”: Application ID

It is important that we keep the order in which a message is logged since that might be important when tracing information.

<1> – Date

<2> – Time

<3> – Port or PID

<4> – Process/Program creating the message. This is used to help track all messages being generated by a specific process. For example, a backup or file resize process.

<5> – Log Message. Since these logs are about tracing what the user is doing, this message may contain anything that the developer deems important. This information is designed to be used for structured error messages. That is what the System logs are for.

<6,n> - Additional Log IDs. This would be like a system log, audit log, and log record ID that would contain more detailed structured information about the Log message. Example: SYSTEM.LOG*5514585-588ASS5-55555

<7,n> - Stack/Trace information. If this is a major error, then there should be tracing information. This information is all about how the process got to the point where we needed to log the error. It may be the subroutine CALL stack, or it may be a programmer's description of the steps that got to this point.

Security Logs

Security logs are exactly what you think they are. Tracking data on security and user access both belong here. They should be persistent, to a point. The key idea is to provide a single place to look for user accesses and overrides. This is where you go to investigate if the user is following policies or doing something else.

Security logs should have the following structure:

ID – Unique ID

<1> – Date

<2> – Time

<3> – Port or PID

<4> – Account/Database ID

<5> – User ID

<6> – Process/Program creating the message. This is used to help track all messages being generated by a specific process. For example, a backup or file resize process.

<7> – Security Log Type: SUCCESS, FAILURE, OVERRIDE SUCCESSFUL, OVERRIDE FAILED LOGOUT, CHANGED PASSWORD, etc.

<8,n> – Security Trace Value Name

<9,n> – Security Trace Value

An example of using the Security Trace Value and Value Name would be on a price override. The information I would put into these fields would be the record that the override was associated with, the amount that was to be overridden, and the new amount that it is becoming.

```
FILENAME] ID] ORIG.AMT] NEW.AMT  
ORDER] 1234] 130] 250
```

Change Logs

Change logs are sometimes confused with audit logs since most people are trying to track what data has changed and when, based on the user. That is half the purpose of a change log. It has to go beyond that and give us what we need to backtrack, and possibly reverse, the changes.

Change logs are usually associated with file triggers, but can be managed by an application directly as well. Change logs should be transitory. They are not required to live for extended periods of time. If you need to keep this information longer, then your developers should be creating an archive file, not a change log.

Change logs should have the following structure:

ID – Unique ID

<1> – Date

<2> – Time

<3> – Port or PID

<4> – Account/Database ID

<5> – User ID

<6> – Process/Program creating the message. This is used to help track all messages begin generated by a specific process.

<7> – Record ID

<8> – File Name

<9,n> – Change Action: UPDATE, INSERT, DELETE

<10,n> – Field Position Changed. AMC or AMC,VMC

<11,n> – Old Value

<12,n> – New Value

<13,n> – Stack/Trace information. This should be the subroutine CALL stack, if available. At least, it should be information on how and where the information changed from.

Logging Notifications

While writing all these files is important, they lose much of their effectiveness if you aren't notified that the log has been updated. Most of what goes into the records can wait but there are times when developers and admins need to be notified when a log is updated.

As part of the overall workflow, a process needs to exist to watch for specific log messages and send that information to a notification processes. Bor-

rowing from the SNMP logging protocol, each log type needs to have a hook (a place where subroutines CALLS can be added). These subroutines would do the work of deciding what details need to be sent, and which notification option should be used.

Some of the notification options that should be available to your application:

- Nightly Reports
- Email
- SMS Text
- IM (Instant Message) Text
- SNMP
- Syslog
- Windows 10 Push Notification
- Mobile App Push Notification

Unique Record ID Management

Since MultiValue databases leave the record ID management up to the software developer, this must be managed by the framework. Other database applications have a few record ID management features built-in due to how the databases are required to store information. In this case, I'm going to assume we will want to manage this information ourselves. Since I am reinventing the wheel, it seems reasonable.

Sequential IDs

The most common type of ID management that most people develop is the sequential ID. This is the easiest way to create a unique ID for every record, but sometimes generating a sequential ID can cause deadlocks that slow down other applications that are also using the sequential IDs.

Pros:

- Always unique
- You will always know exactly what the next ID will be

Cons:

- Record locks (deadlock) can occur if two processes are trying to increment the ID at the same time
- Additional disk processing when read/writing the sequential counter

GUID/UUID

GUID and UUID are basically the same. The only difference is that UUID is the open source way of creating a unique ID, and GUID is the Microsoft way. They both do the same thing, making sure the ID is unique based on the machine you are creating the ID on.

If you have never seen a UUID, it looks like the following:

30dd879c-ee2f-11db-8314-

0800200c9a66

GUID are actually really large numbers, but they are represented in Hex to make them easier to read. GUID and UUID are calculated, so there is no reason to store a counter anywhere, but they are unique across applications and often even across hardware platforms. This makes them ideal for creating IDs when you have a federated application with a central office holding all the information, and remote store or offices holding only their information.

Pros:

- Always Unique. Even across different databases
- Calculated Value; so no additional disk overhead or locks

Cons:

- Calculated Value; May take up more processing depending on how the ID is generated
- Won't know what the next ID is, as it is always unique

Date/Time IDs

Date/time IDs are great for logs but do sometimes require additional processing to make sure they are unique. Date/time IDs are exactly what is being described; an ID that is made up of a the current date and time.

Many times these IDs also include Port or PID numbers to offer additional uniqueness, and to prevent record collisions.

A common structure is:

DATE*TIME*PID

The main drawback to date/time IDs is the millisecond factor. You need to include something to handle the mil-

lisecond factor so your applications don't have a record ID collision.

Most of the MultiValue Databases have a way to handle the millisecond factor, but if your system doesn't, then you need to do the additional tests to make sure that the ID doesn't already exist on file.

Pros:

- Mostly unique, but has issues with millisecond factors
- Calculated value; so no additional disk overhead or locks

Cons:

- Additional checks may be required so you don't have record ID collisions
- Must handle the millisecond factor, and must include a PID to keep uniqueness

Conclusions

While I've only talked about logging and record ID management, they are key framework components that will be used throughout the rest of the project. Logging is not something that should be left for last, but part of the core framework.

Stay tuned for the next part of this series in the next issue. **IS**



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

Is Your Membership Profile Up to Date?

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

Please go online and update your membership profile today!

intl-spectrum.com/membership