


\$7.00 U.S.

INTERNATIONAL SPECTRUM®

THE MULTIVALUE  TECHNOLOGY MAGAZINE | JULY/AUGUST 2017



MULTIVALUE BENCHMARKING

ALSO IN THIS ISSUE:

- open "makeover"
- JSON vs. XML: Fight!
- Business Tech: UX/UI Part II



INTERNATIONAL SPECTRUM

37TH ANNUAL MULTIVALUE CONFERENCE

DESTIN
FLORIDA

The Premier MultiValue User Event

APRIL 23 – 26, 2018

THE HENDERSON
DESTIN, FLORIDA

www.intl-spectrum.com/conference

The **International Spectrum MultiValue Conference** is the only place that provides diverse solutions and support for integrating and developing your Enterprise applications into the world of business and technology.

This annual event is the place to find the **knowledge, people, resources,** and **technologies** for enhancing your MultiValue (D3, UniVerse, UniData, Reality, jBase, Caché, QM, and OpenInsight) software applications.

Save the Date!





COVER

6 **MultiValue Database and Framework Benchmarking**

Benchmarking has been a topic that's been around for several years, but nothing really has been provided to the MultiValue Community to use. I've heard many different reasons for the lack of benchmarking, so I'll go through a few of them. **BY NATHAN RECTOR**

FEATURES

11 **The Six Million Dollar Open** The term extensible means that a language or format can be tweaked, adjusted, or expanded by the users. It can go beyond what the original designers had in mind. MvBASIC is extensible. The open command is an excellent candidate for an extensible makeover. **BY CHARLES BAROUCH**

16 **Expanding Your Toolkit: JSON vs. XML** MultiValue strings are an excellent way to express data compactly. In a world where, increasingly, every system seems to need to talk to every other system, it is useful — and desirable — to have as many communication formats available to you as is reasonable. Last time out, Bennett Barouch gave you a good look at JSON. This time he's pulling back the curtain on XML. **BY BENNETT BAROUCH**

19 **UX/UI Part II: Assumptions** Interfaces work, or fail, based on agreed reality. The challenge isn't solely about languages. It has to do with culture, generational changes, and real world experience. Computer programming is a dialog between the programmer and the user. Let's start questioning the assumptions we make in creating our conversations. **BY CHARLES BAROUCH**

DEPARTMENTS

From the Inside page 4

From the Press Room page 13

From the Inside

We're already planning for the next Spectrum conference. If you haven't yet seen the information for the 2018 Spectrum, it will be April 23rd-26th at The Henderson in Destin, Florida. This is a new venue for the conference.

Looking at a calendar, it seems like the conference is pretty far away, but it really isn't. We are currently working on the planning for sessions, so you have something to show to your company. The Spectrum Conferences are designed to help you in your business. That means that we want to know what you, as a MultiValue user, would like to see. What topics will sell your bosses on sending you to Spectrum?

Some Thoughts

The demands on your IT department are getting more specific, which defines the types of examples and business solutions that you need to know more about. I remember having to justify my attendance, so I watch the trends that are coming. The conference works because we update our content with the times.

We have MultiValue Experts as speakers who give us great content, but it is up to me to guide them. It is up to you to guide me. I would like to hear from the MultiValue Users and Developers. I want to make sure you get the materials,

education, and experiences that are needed to maximize your ability to create efficient enterprise applications.

If you are unsure what you would like to see, here are a few that I've noticed:

Virtualization/Cloud Computing

Everyone is talking Cloud Computing and Virtualization, and for pretty good reasons. We have presented topics on Virtualization and Cloud Computing at the conference before, but the topic is deep. What are you interested in hearing next? Do you need to know how to integrate MultiValue databases into a cloud API, like Azure or Google App Engine? Or are you looking run a MultiValue server in a Virtual Machine? Maybe the plan is bigger than just the database server?

Tablets and Mobile Apps

Mobile Line-of-Business applications are a must. Many companies are not taking full advantage of what mobile apps can bring or how they should be used in the enterprise. How can we help you help them?

Windows 10 and Universal Apps

Working with Windows 10 development can be a pain. The new Universal Apps and Windows 10 features are nice, but very different from what most developers are used to.

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 2017 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.

Would you like to see general sessions on how to use Windows 10? How Universal Apps are created? These would probably not be MultiValue specific, but we need to know about other software to make our jobs possible, so I'd be happy to include them.

Office 365 Integration

Looking to integrate Office 365 more closely with your Line-of-Business application? Documents, Calendars, Email, Spreadsheets... there are versions that are free and versions for a fee. All can be integrated with your MultiValue data.

Other Integration: Amazon, SalesTax, Email, CRM, EMR, LDAP, Active Directory

Do you have other applications within your environment that you need to integrate with? We have been an island unto ourselves for so long, we sometimes forget that our business has information that other systems can use.

What we have in our system is increasingly necessary for other business functions. What integration options do you need?

External Programming Languages

What languages are your business requiring you to work with? .NET, Java, Python, Javascript, HTML, JSON... there are many others. We can provide topics on all of these, both stand-alone and in connection with your databases.

These are just a few examples. What do you find most interesting? What trends and topics that you would like to have covered in sessions at Spectrum? I would like to hear from you. Contact me at nathan@intl-spectrum.com. **IS**



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

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



taking multivalued ...
where it has never been before

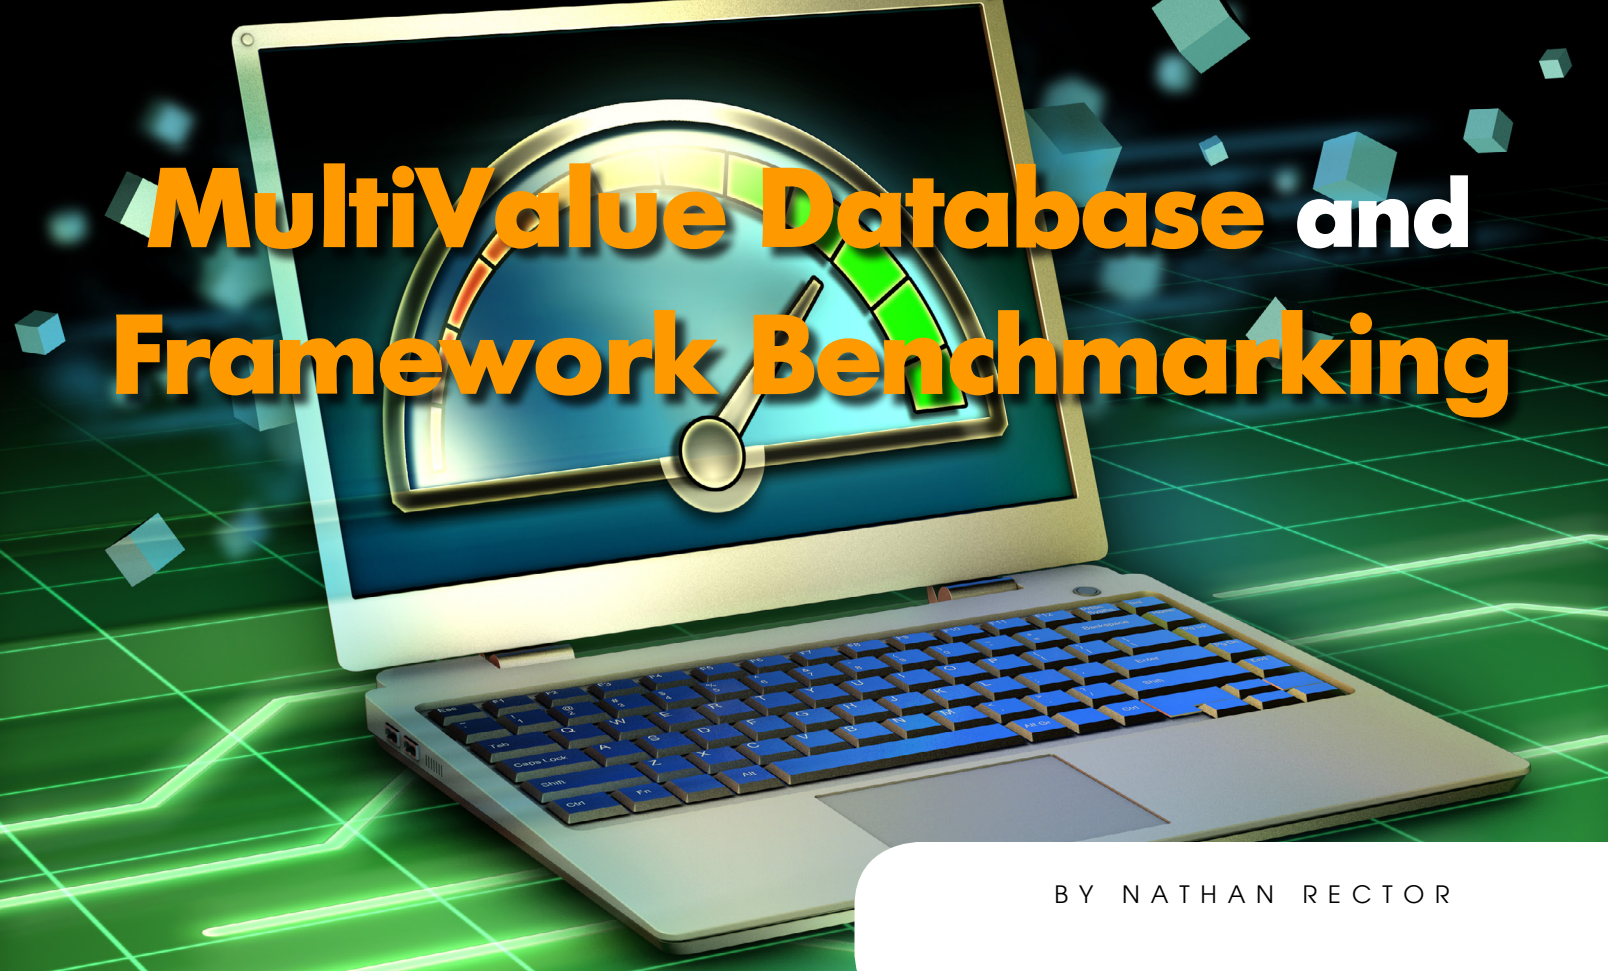
Ladybridge Systems Ltd

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

Worldwide distributor: Zumasys, 9245 Reasearch Drive, Irvine CA 92618, USA
www.zumasys.com

- Close compatibility with most other multivalued 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



MultiValue Database and Framework Benchmarking

BY NATHAN RECTOR

Benchmarking has been a topic that's been floated around for several years, but the MultiValue Community doesn't have much in the way of prepackaged tools. I've heard many different reasons for the lack of standard benchmarking tools, so I'll go through a few of them.

What are TPC.org Benchmarks?

There are many benchmarks out there for a wide variety of other environments, but I'm going to focus on the most popular. For many years, the TPC.org benchmarks have been what most people think of when talking about Benchmarking. You've probably heard of the TPC-C, TPC-H, and TPC-E versions. These have become the industry standards as benchmarks that people want to use. But, like any benchmark tool, they all have drawbacks as well.

Hardware keeps getting faster and well-made software wants to leverage those advances.

The TPC-C benchmark is the most common, mainly due to how old it is. TPC-C has been around since 1992. It is designed to model an OnLine Transaction Process (OLTP) wholesale application. Since it was designed over twenty years ago, it made some assumptions about the hardware that don't hold up today. Mostly, it expected that machines did not have a lot of RAM and it expected inefficient Disk IO. The TPC-E benchmark was created as a more modern replacement for the aging TPC-C benchmark. TPC-E uses an OLTP stock exchange as its data model. The design is more sophisticated in its simulation of a real-world

system. It takes into account the fact that modern servers use lots of RAM and other system resources to offset the slower Disk IO.

The TPC-H benchmark is designed to model ad hoc reporting. Its main purpose is to measure query/select results using a Data Warehouse as a database model.

Keep in mind, TPC.org has many more benchmarks that provide information for different purposes, but these three are the most common. While the TPC council offers all the details needed to run these benchmarks, you are not allowed to release your benchmark test to the general public without getting an official sign-off from them.

Is Benchmarking Dead?

Benchmarking used to be the central marketing piece of comparing databases, but has fallen to the wayside in

recent years. Many C-Level managers still look at that component without realizing it, but there's a reason benchmarking isn't as critical nowadays: Hardware.

When it used to cost the IT department their first born in order to throw hardware at a problem, speed was more about the software. Now, RAM costs you pennies, and readily-available Solid State Drive (SSD) elevate Disk IO to the speed of RAM. On top of all of that, you can now create clustered systems which distribute the database processing across more than one machine.

Virtualization also adds a wrinkle into the mix because a Virtual Machine (VM) can look like a single machine, but have its physical resources allocated much differently to improve performance. And you can often tweak the VM's hardware by changing a few software settings, instead of physically replacing hardware components. Cloud-based virtualization often doesn't even require the machine to be rebooted for these hardware changes to be made.

Of course, the operating systems available to the databases have changed as well. A benchmark running on a Windows platform will generate different results than one running on Linux, even on the same hardware configuration running the same database version. You can even get different results running the same operating system and the same hardware, but tuning the operating system's kernel differently.

Now add the next variable: network infrastructure. As we know, in the real world, our network infrastructure affects how fast data can be sent to our servers. We no longer have dedicated connections to the database like in the old RS-232 days. Network issues, such

as client connections, switch overhead, and network traffic, all will affect how your benchmark produces its information.

After addressing the hardware configuration, infrastructure, and the operating system variables, you still have to look at the database version on which you are running the benchmark. Each release has gone through improvements in order to take advantage of modern server architecture and resources. Hardware keeps getting faster, and well-made software wants to leverage those advances.

There has also been improvement in database system internals, like indexing, caching, and other efficiencies, that the benchmarks don't know about as well. For example, the database may use an inverse index in place of a b-tree, or even a bitmap index, based on the type of data being stored. Or it may choose to load the whole file table into memory and keep it there instead of relying on Disk IO. These change the whole CRUD (Create, Read, Update, Delete) dynamics.

Now that you have considered all of the above, there is one last variable to consider: administrator-introduced performance features. When he or she places an additional index on a secondary key — or uses a stored-procedure (BASIC routine) which utilizes a database specific feature — the results can become skewed even more.

All these things have changed the simple "Benchmark Marketing" concept of "I'm faster than you are" that used to go on in the early days of computing.

Company:	HP
System:	HP ProLiant DL380 G7
OS:	Windows Server 2008 R2 Enterprise
Database:	Microsoft SQL 2005 Enterprise SP3

Fig. 1

What Does This Mean?

Are benchmarks dead? Are they no longer useful? No, benchmarking still has its place, but it ceases to be the apples-to-apples comparison it used to be. Marketing used benchmarks when it was reasonable to assume that the machine with the faster benchmark would be faster in production. Now it is hard to assure that a specific instance of a system will produce the same benchmark values twice, much less assume that a generic benchmark is an indicator of how your system will react to the daily workload.

Benchmark comparisons have moved from the comparison of different databases to the performance of the hardware that database is installed on. If benchmarks are used provide performance numbers for hardware configurations instead of database comparisons, you will find the value in them.

You should also be very skeptical about published benchmarks you may find. There is a lot of important information that is missing from them. Let's consider two benchmarks from HP, published on TPC.org, for the TPC-C Benchmark.

Both start with identical specifications [Figure 1].

But the reported benchmark results are not the same [Figure. 2]

Looking at this you might be scratching your head. The specs seem to say

Submitted	Performance
05/04/11	1,024,380tpmC
05/11/10	803,068tpmC

Fig. 2

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



they are the same system, right? You need to dig into the details to find the differences [Figure 3].

Looking at this information, they *are* pretty close, other than the amount of memory and the total number of clients used. A natural assumption is “oh, they just threw more memory at the problem,” but let’s look at something else they did that really makes a difference [Figure 4].

First thing to point out is that they spread the OS, log files, and database storage onto three different disk volumes. Not really a bad thing and should be done as general practice for most database servers. But this will affect your comparisons if you currently don’t have your systems setup this way.

The next thing to point out is that the database storage is on SSDs, and not on spinning disks. While they are both on SSDs, which is nice, you will notice that the quantity of SSDs are different. That means the database is being split across more disks differently, and the information sheet provided does not state if there is a RAID setup, how the RAID is setup, or how the data is allocated across the disks. Regardless, more disks mean more seeks per second when the database looks for data to satisfy a query.

As you can see, with only a slight change — from a cost standpoint — in hardware, you can get very different performance numbers. The other thing to note in the performance numbers that is very telling is the “Number of Clients.” This variable matters, especially with TPC-C benchmark.

As you probably have guessed, the number of clients tells you how many connections you have to the server. Also, the ‘tpmC’ value is an accumulative average based on all the clients working together. That’s right, the more clients working on a system, the better the performance number looks.

If you do some simple math, taking the performance number and dividing it by the number of clients, you will

notice that the performance per client is really about the same [Figure 5].

The published benchmarks look good, but they are misleading. Not only do hardware changes affect performance numbers, you have to consider how the benchmark was run, such as number of clients and how long the benchmark was run for, to get a good idea of true numbers.

This also means it is really hard to compare numbers between vendors and databases because each one is run differently. Many vendors and/or database providers have their own implementation of the benchmark software as well, so there can be major variations between performance number based on how, when, and how long a benchmark was run.

Using Benchmarks in MultiValue

Another major drawback to benchmarks is the fact that they are based around using First Normal Form (FNF) and SQL syntax. This isn’t really a big deal with MultiValue platforms since we can do both, as well as non-SQL, but it also means you can’t use standard benchmarking tools either, because they can’t test all of the ways we access data.

Performance:	1,024,380tpmC	803,068tpmC
Number of Clients:	839,690	639,960
Memory:	384GB	16GB
CPU:	Xeon 3.4Gh	Xeon 3.3Gh
	2 Processor/6 cores	2 Processor/ 8 Cores
	12 threads	8 threads

Fig. 3

Controller:	8gb Fiber	8GB Dual Channel
Disk (Data Storage):	81@ SSDs	128@ SSDs
Disk (Logs Only):	12@ SAS 7.2K	54@ SAS 7.2k
Disk (OS):	2@ SAS 15K	2@ SAS 15k

Fig.4

Performance:	1,024,380tpmC	803,068tpmC
Number of Clients:	839,690	639,960
Per Client:	1.21	1.25

Fig. 5

So where does that leave us? Doing it ourselves.

As a proof of concept to see what it would take, I decided to create a TPC-C-like benchmark. I chose this benchmark over TPC-E and TPC-H because it was the easiest. There are more examples floating around for how to implement TPC-C than there are for TPC-E.

In addition, TPC-E requires the use of C code provided by TPC.org, so I didn't want to try to implement or translate that code into MultiValue BASIC or some other client. Several MultiValue versions play well with C and I invite you to implement TPC-E if you are so inclined. If you do, please send me a copy so it can be shared. Another thing to consider about TPC-E is that the only vendor that has done any work with it is Microsoft. While it was designed by Microsoft, SAP, and Oracle, only Microsoft has jumped on the bandwagon to officially produce benchmarks for it.

TPC-H requires the use of SQL queries. While I could translate those SQL queries into MultiValue Retrieval statements, or use an SQL parser, it required a lot more setup and database configuration than I had time for since this is just a proof of concept.

With all that in mind, I settled on the TPC-C benchmark. This has the added benefit of mimicking the model of many of our day-to-day systems

anyway. This benchmark also fits the Port/Pib/Telnet connection style that many of our MultiValue applications currently use.

The last benefit, since TPC-C is also a Disk IO intense benchmark, is that it tests a key advantage of the MultiValue Database Framework.

The TPC-C programs I wrote are in pure MultiValue BASIC, and are designed to adhere as closely to the TPC-C spec as possible. Since each client runs on its own port, just like any other user would, my implementation follows the TPC-C spec much closer than most TPC-C benchmark programs do. In order to decrease the cost of benchmarking, it not uncommon for vendors to use a multi-threaded client on a single PC, instead of one PC per connection, in order to mimic the one-client connection that the TPC-C benchmarks states.

Since my programs are designed be as classic to the benchmark as possible, it does not take into account any performance improvements which are commonly used in our day to day systems. For example: I did not implement b-tree indices or version-specific SELECT statements. This version of the benchmark does not take into account any data structure optimization that could be done, either. All the test data is still stored in FNF, even though a huge speed improvement would be possible by changing it to use MultiValues and nested tables.

ter – they certainly are for this sort of test – but to try to approximate how a MultiValue database would handle data when limited to the same format as an SQL style database.

Here is my test system and what I did [Figure 6].

Now, I would love to scale this up to sixty thousand or ninety thousand clients, but I don't have a server or the licenses to run that kind of test. If we were to scale this up statistically, even with say a forty percent decrease in performance due to increased disk IO and memory, that would still put the performance at 6.6 million tpm. And that's the plain-Jane, done-in-a-rush, not-very-optimized version.

As I've stated above, there are many different variables that would affect these numbers, so it really isn't fair to compare performance numbers from one database type to another; or to compare a purely statistical number to an actual number.

The point is that MultiValue systems are very speed-competitive with other systems; not slow or behind-the-curve. This exercise was focused on recreating a standard benchmark. It was very informative seeing the kind of results I could generate. There are still many different variations on this benchmark. I would be very interested in people's feedback. And, if you happen to have a big system to run this version on, or want it as a jumping off point for your own, more tuned-up version, let me know. **IS**

System:	Dell PowerEdge T610
OS:	VMware/RedHat EL3
Database:	D3/Linux
Disk:	SSD
RAM:	80
Duration:	2 hours
Performance:	1179tpm
Number of Clients:	10
Per Client:	117.9

Fig. 6

Again, the idea wasn't to optimize the benchmark to prove that MultiValue data structures are bet-



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



GET CONNECTED.

KNOWLEDGE AND EDUCATION FOR THE MULTIVALUE PROFESSIONAL.

ABOUT OUR PROFESSIONAL MEMBERSHIP

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

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

Professional Membership Includes:

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



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

Open

BY CHARLES BAROUCH

In the opening to the TV show, Six Million Dollar Man, they promised to rebuild him – better, faster, and stronger. This occasional series, the Six Million Dollar _____ will offer suggestions on how to elevate mvBASIC commands by making them better, faster, and stronger. We have the technology.

Additionally, the common area, combined with the equates, will keep the name of the file handle consistent through all of your programs.

```
BP OPENER
001 subroutine OPENER(D.LEVEL,FNAME,MODE)
002 * by Charles Barouch (Results@HDWP.com)
003 * as seen in International Spectrum Magazine
004 include BP FILECORE.INC
005 if unassigned(FILE.STATUS) then FILE.STATUS = ''
006 *
007 OP.FLAG = 1;* Yes, we should open the file
008 if D.LEVEL = 'DICT' then
009     * Don't open it if it is already open
010     locate('DICT-':FNAME,FILE.STATUS,1;FS.POS) then
011         if FILE.STATUS<2,FS.POS> # '' then OP.FLAG = 0
012     end
013 end else
014     locate(FNAME,FILE.STATUS,1;FS.POS) then
015         if FILE.STATUS<2,FS.POS> # '' then OP.FLAG = 0
016     end
017 end
018 if MODE = 'RESET' or (OP.FLAG) then
019     open D.LEVEL,FNAME to FILES(FS.POS) then
020         FILE.STATUS<2,FS.POS> = 1
021     end else
022     call CALL.AUDIT('OPENER',FNAME:' did not open',' SHOW,ABORT')
023 end
024 return
```

Fig. 1

```
BP FILECORE.INC
001 * by Charles Barouch (Results@HDWP.com)
002 * as seen in International Spectrum Magazine
003 common /FILECORE/ FILES(200), FILE.STATUS
004 *
005 equ DICT.ACCOUNTS.RECEIVING.FILE to FILES(1)
006 equ ACCOUNTS.RECEIVING.FILE to FILES(2)
```

Fig. 2

Cold Open

Consider the OPEN command:

```
OPEN 'ACCOUNTS.RECEIVING' TO B
ELSE NULL
```

I hope I don't have to point out why opening a file to a single letter variable is bad. This example gets extra demerits for using a letter that doesn't even appear in the table name.

Likewise, using NULL in the ELSE clause will suppress all error reporting. While there are edge cases for doing that, generally we want to know when an OPEN fails. Unfortunately, this command will compile, making your life harder.

Since we can't rebuild and replace the actual OPEN command, our best course is to wrap the command in code which makes it more robust and consistent. Have a look at this alternative [Figure. 1 and Figure. 2].

Why is This Better?

Any errors generated by any OPEN command will be logged the same way. This will make audits (internal

PICK PROFESSIONAL

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

CALL THE BEST! Contact...

Matt Hart

EXECU-SYS, LTD

1411 Broadway, Suite 1220

New York, NY 10018

(212) 967-0505

(800) 423-1964 x 302

Fax: (212)947-8593

Email: mh@eslly.com

Consultants Needed Nationwide

and external) much easier and improve your ability to fix your system. As a bonus, the common area will keep the file open for the lifetime of the user's session, making future requests to open faster because they will only be confirming status, not doing a real open. Additionally, the common area, combined with the equates, will keep the name of the file handle consistent through all of your programs. That's a huge win when debugging and documenting.

While those three benefits are significant, there's a bigger win. If you want more information, like the name of the calling program, the port number, the time and date, or anything else, you can add those features to this one routine and you gain that information in every program that previously would have had to perform an open sepa-

rately. You could also make your version of the OPENER program smart enough to auto-create workfiles when they are needed. **IS**

Got your own wrapper for a MultiValue BASIC command or TCL verb? Send us your Six Million Dollar make-overs so we can share them with the community. Email editor@intl-spectrum.com.



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, Pax Solaria, PereheliionSF, and the Interrogative series, which begins with Tiago and the Masterless.

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



FROM THE PRESS ROOM



Kosday Solutions Releases Excel Add-in for MultiValue

Kosday Solutions releases Excel Add-in for D3 and QM. With Kosday's Add-ins for MV databases, the connection and following exploitation can be carried out in an easy way. Besides the power of queries running as complex as wished, you can quickly gain access to concrete data through the use of available Excel functions. It has never been so easy to customize documents or produce reports. ■



Entrinsik Informer Named a 2017 Business Intelligence Top Performer in Customer Success Report

Entrinsik Inc. has recently announced it has been named a "Top Performer" business intelligence tool in the "2017 Summer Customer Success Report" published by FeaturedCustomers.com. FeaturedCustomers.com is the number one platform for find-

ing and reviewing IT software through verified customer success content. With nearly 200,000 customer endorsements from some of the most recognized brands, FeaturedCustomers.com helps its prospective users make informed IT purchasing decisions.

Entrinsik Informer simplifies the often complex process of accessing, cleansing, blending and analyzing data from multiple sources with its flexible, extensible architecture and intuitive user interface. Informer enables organizations around the world to extract critical insights buried in data, whether in legacy systems, spreadsheets, or in the cloud.

The Customer Success Report is based on the curation of authenticated case studies, testimonials, and reviews from across the web, telling a complete story of vendors in the BI space through real customer experiences. Scores are determined through a number of factors including number of reviews, customer success content, social media presence, and vendor momentum based on web traffic and search trends.

The Top Performer designation is awarded to vendors with significant market presence and resources and enough customer success content to validate their vision. Highly rated by customers, Entrinsik Informer consistently publishes customer-driven content and curates dozens of positive reviews from multiple third-party sites.

"Entrinsik is dedicated to our customers' success. We are always honored to be recognized for this dedication and are excited to be included in FeaturedCustomer.com's report. It goes to show that when a company listens to its customers, their products can only evolve to be better and better," says Doug Leupen, CEO at Entrinsik.

Entrinsik Informer's profile, curated content, and the full Customer Success Report can be found at <https://www.featuredcustomers.com/vendor/Entrinsik-informer>. ■



Ladybridge Systems Releases QM 3.4-11

Some more new features requested by our users...

The locking system now records the time at which a record or file lock was acquired as an aid to debugging locking issues. The time is included in the LIST.READU report if the DETAIL option is used and also appears in process dump files or deadlock diagnostics.

The optional replacement character in an MCP conversion code can be enclosed in single or double quotes. Empty quotes cause non-printing characters to be removed without replacement.

The I-type compiler now supports a subset of the options

to the \$MODE QMBasic compiler directive.

Overwriting an existing local catalogue pointer in the VOC now queries for confirmation if the pathname of the object code is different.

The Hot Spot Monitor (HSM) now reports i/o counters (open, read, write, delete).

A new @-variable, @PSTAT, has been added. This can be used to set diagnostic information that will be displayed by the PSTAT command.

The USER, GROUP and MODE keywords have been added to CREATE.ACCOUNT, CREATE.FILE and CONFIGURE.FILE to allow setting of ownership and access permissions on Linux/Unix systems.

The SAVE.LIST, GET.LIST and DELETE.LIST commands now all support an alternative file name.

The SHOW command has a new keyword, SORT.SUP to suppress the default @ID sort.

The WRITE.DELETE.THEN.ELSE setting of the \$MODE compiler directive has been extended to include the MATWRITE and MATWRITEU statements.

The SETPTR and SET.QUEUE commands now accept the AT clause in mode 5 to identify the printer to be used on the client PC. This is applicable only to AccuTerm when using a terminal type with the -at suffix.

The FL\$LOCKS key to the QMBasic FILEINFO() function

FROM THE PRESS ROOM

returns a list of the locks in the file owned by the calling process. Attempting to enter the debugger in a phantom process raises exception SYS.PROGRAM.NO_PDEBUG if no interactive process is waiting in the PDEBUG command.

The QMClient API now includes QMSystem() to return the value of a SYSTEM() function on the server and QMTrans()/QMRTrans() to perform the equivalent QMBasicTRANS() and RTRANS() operations.

Setting the DEADLOCK configuration parameter to 2 extends the action of mode 1 to save the state of the lock tables as an aid to diagnosis. ■



MV# Developer Studio Release 2.0 Now Available

MV# Developer Studio from ONgroup Intl is an editor for MV BASIC and an IDE for MVON# Netbuilder applications. Netbuilder runs applications in .NET with SQL Server that were originally written in System Builder (SB+). The Studio includes a designer, property pane and other features for rapid development and maintenance of applications.

MV# Developer Studio 2.0 has been enhanced to be a contemporary editor running in .NET that works for all

MVON# customers. MVON# Netbuilder sites still get the big benefits of an IDE, and now those running MV BASIC and related languages without Netbuilder can still get the benefit of this powerful MV# editor. Compiling with the MV# Developer Studio transpiles to C# and runs the MV code in .NET.

Included in the feature set are color coding, code folding, and multiple tabs, including across servers and databases. "Having been an ardent fan of Easy Developer (ED)," says Peter Carroll, Senior Software Engineer at ONgroup Intl, "I have to admit that using the Studio is a big advance for application development. I highly recommend contacting ONgroup and giving it a spin."

Clicking or touching the screen to compile, with or without debug, is what developers working in other languages expect. "It is one place in the demos to MultiValue developers where we hear 'oohs' and 'ahs,' however," said Grant Hart, Chief Software Architect for ONgroup. MV# Developer Studio 2.0 feels more like Microsoft Visual Studio than other editors that support MV BASIC, PROCs and Paragraphs.

The MV# Developer Studio is part of the MVON# suite of tools that free MultiValue applications from using MV-specific databases or MV-specific run machines. MVON# tools run MultiValue applications in .NET with MV data in SQL Server. ■



Kourier Integrator Release 4.4 Introduces API Subscriptions and Swagger Documentation

Kore Technologies announced that Release 4.4 of its award-winning Kourier Integrator Data Management Suite is now generally available.

The primary focus for this release is to enhance Kourier Integrator's support for RESTful Web Services with the addition of API Subscriptions and Swagger API documentation. These advanced features support Kourier's mission of being the most complete and easiest to use real-time integration solution for MultiValue systems.

API Subscriptions strengthen Kourier's enterprise application integration capabilities by providing a new method for communication between disparate systems via the Webhooks paradigm. Kourier reduces the time and effort required to setup and implement Subscriptions for the MultiValue API developer and the API consumer. By subscribing to published APIs, integrating applications automatically receive notifications in real-time when there

are changes to the data in the MultiValue system.

The addition of Swagger (OpenAPI) documentation to Kourier REST is also significant. As the emerging standard for REST documentation, it makes both API developers and consumers more productive. Developers can automatically generate interactive documentation from Kourier's REST Resources to help them visualize and test the interface during development. API consumers benefit from being able to quickly deploy integrations by converting Swagger API definitions into code for over 40 programming languages. ■



Zumasys Introduces MV Connect

Zumasys introduces MV Connect for companies that rely on access to web, mobile application, and cloud solutions. MV Connect creates modern, browser-based user interfaces that are built on Web 2.0 frameworks, such as Angular JS. It will allow integration with popular SaaS-based applications like Salesforce™, Workday™, and NetSuite™.

MV Connect creates a RESTful web service entry point for UniData, UniVerse, D3, or OpenQM. This opens access for mainstream developers using familiar tools and still

FROM THE PRESS ROOM

maintain, enhance, and use decades of priceless code. ■



Rocket Software Continues Modernization of UniData with Python Support

Rocket Software has augmented the functionality of its award-winning UniData MultiValue Database, part of the MultiValue Application Platform, with the addition of Python programming support and audit logging capabilities. These enhancements ex-

pand the potential user base for UniData and improve the ability to recruit new development talent, and allow users to easily establish configurable histories of interactions, events, and activities.

“Being involved with UniData for more than 30 years and still firmly believing in the rock-solid 24/7 performance of the database, we now have a breath of fresh air blowing in the form of Python for U2,” said Jan van Schalkwyk of Australian insurer YOUI. “This much-anticipated release will extend the life of many applications way into the future and beyond. Modern applications can now be developed with a wealth of resources available.”

The inclusion of Python support opens up MultiValue databases to architects and programmers who use one of the world’s most popular languages, helping organizations that rely on MultiValue to expand their hiring base.

In addition, because companies must conform to audits in multiple formats and from various organizations, audit logging lets UniData users create audit histories to help meet standards such as HIPAA, PCI-DSS, and the upcoming European Union General Data Protection Regulation (GDPR). UniData now makes spot audits easy with more granular audit data and access to chronologically generated data, and the new software includes support for

sequential file logs, creating improved performance without system interruptions.

“After performing rigorous testing of this new version of UniData, the features and flexibility of this product answer almost every conceivable audit question,” said Russell Patterson, IT Specialist for Rural Finance, a specialized lender fostering sustainable growth in Australia. “The Audit module can be configured from high-level system-wide activity to specific user file-level access and change logging, to program and process logging. UniData Audit has all the answers. If you can think of an audit question, UniData Audit can be configured to provide the answer.” ■

eXtra Newsletter

Stay on top of Industry News

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



www.intl-spectrum.com/newsletter

JSON vs. XML

BY BENNETT BAROUCH

Back in the stone age — the early 1980's — SGML was created to express document structure, so that complex documents could be electronically shared and rendered by cooperating government agencies and companies. “This part of the content is the title,” “this part is a figure,” “this is a legend under a figure,” and so on.

Still in the stone age — in the 1990's — HTML arose to express less complex documents in a web-browser context. At the dawn of the modern era — the late 1990's — XML was created with the intent of finding a middle ground between the two. One that was both human-readable (plain text with Unicode support) and machine readable (sufficiently structured to be efficiently and unambiguously parsed). It was intended to be a practical compromise between completeness and simplicity.

Like its parents, XML was conceived as a means of expressing a *document*. However, it was understood and desired that it could be used to express non-document data as well. XML's elegance comes from not having a rigidly fixed lexicon. It's all well and good to say `<title>` is a tile, and `<p>` is a paragraph, but how cool is it to say that

It's all well and good to say `<title>` is a tile, and `<p>` is a paragraph, but how cool is it to say that `<something>` is a something?

`<something>` is a something? Without having to define an unworkably large universe of possibilities in advance and then deal with missing features, and force-fit messes they cause us to create, we can suddenly talk about `<inventoryControlNumber>` and `<favoriteIceCream>` just as easily as anything universally pre-defined.

If the call-a-thing-a-thing philosophy reminds you of JSON, it's not an accident. JSON was designed to be *XML-light*. We get to decide what terminology best suits the task at hand in JSON because that was one of the aspects that was kept from XML. This

of course begs the question: So why do we need both XML and JSON? Good question.

The JSON and XML snippets in [Figure 1] convey the same core payload.

You can see that JSON is a bit more compact without being any more obscure, so that looks like a win. In situations represented by the above example, it is. A huge set of real-life use cases are well served by JSON. It's terrific. No surprise that it is the data packaging system most in use on the web to-

JSON:

```
{
  "employees": [
    {
      "firstName" : "Albert",
      "lastName"  : "Einstein"
    },
    {
      "firstName" : "Albert",
      "lastName"  : "Schweitzer"
    }
  ]
}
```

XML:

```
<employees>
  <employee>
    <firstName>Albert</firstName>
    <lastName>Einstein</lastName>
  </employee>
  <employee>
    <firstName>Albert</firstName>
    <lastName>Schweitzer</lastName>
  </employee>
</employees>
```

Fig. 1

day. You might easily spend your entire career not needing anything more.

But.

For binary values — like images, sound, and video — JSON requires a to-string encoding and from-string decoding. Although easy enough to pull off for an occasional bit of multimedia content, when you have to scale this, it can add up to being expensive computationally, bloated, and there are transmission resources required.

XML can handle binary data expressed as text (just like JSON), or referred to by URL or other external access methods (just like JSON). But it can also manage your multimedia content as an embedded/in-line/raw/un-expanded binary data by employing a CDATA tag. Also, there is something called Xpath, **which** traverses the more com-

plex trees common to XML for you, so you can find things like “/book/published/year/text()” without having to write any custom code. And, XML supports name space segregation, metadata (attributes), and externally or internally defined schema to describe how a payload is to be understood. Hmm. Getting interesting for more complex data scenarios, isn't it?

Oh, and you can XSL to filter and transform one XML “document” into another one more suited to a particular task.

And...

You must get the picture. Applications moving binary data or complex data hierarchies are usually better served by XML. When those requirements are not present, JSON is likely the better pick, due to its compactness, simplic-

ity, familiarity, and widespread use by many of the web services you may interact with.

This brief article cannot provide a primer on XML. For those not familiar with it, I merely hope to have alerted you the fact that JSON has a more powerful ally able to do the heavy lifting that lies beyond JSON's reach. When your work calls for custom packaging of binary and or deeply hierarchical data, a solution leveraging XML will most likely be better than either JSON or a completely proprietary mechanism. **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.

**Celebrating 20 Years
as a leader in the
MultiValue Industry**



D3 UniVerse UniData PICK jBase mvBase Caché

IT Solutions. Proven Results.

One Project. One Decision. One Keystroke at a time.

- Custom developed solutions based on your needs
- Senior level developers and business analysts to guide you
- Developing long term partnerships


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

MARKETPLACE

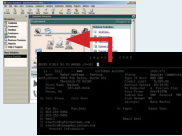
ACCOUNTING

Natec Systems

www.natecsystems.com | nrector@natecsystems.com



- 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[®]
by Zumasys

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

UX/UI Part II

Assumptions

BY CHARLES BAROUCH

In 1984, it was clever to make the Save icon look like a floppy disk. There was an immediate visual connection for any user who was even a little tech savvy. It was easy to explain to the less sophisticated user: “Before you click this disk thing, you need to put one of those disk things in the slot.” Today, the average child is already on the same level as yesterday’s moderate-level users. And, of course, they have no idea what a floppy is, unless they have one of those parents who blew the dust off of one and showed it to them. It won’t be long before the CD and DVD disappear from the collective consciousness in the exact same way.

We live in a world where an adult can conduct their entire life without knowing how to read an analog clock. They’ve never dialed a phone, the don’t know why the soldier in the movie says, “Watch your six” when he means, “Look behind you.” And we are still using iconography that assumes that everyone knows what a landline handset looks like. UX (User eXperience) is about making things easy and obvious to the person using the software (and hardware). That is, except when it isn’t.

Game On

Games, one of the most UX-driven things we do with consumer technology, revel in making some parts of the interface easy and some hard. Players (users) have to earn some of the but-

tons instead of having access to the entire interface. For example: “Level Up! You can now tap the Green Door to access the PvP arena.” In games, developers use the UI against the user to make the UX better for the user.




Fig. 1

Of course, commercial software is different. That is, except when it isn't. Educational software, including professional training software, often requires you to earn access to additional features by passing a quiz or performing some sort of function. We routinely lock users out of features until they enter a password, or complete two-factor authentication, or prove they aren't a robot. When I'm testing web pages, I get asked if I'm human twenty or thirty times in a morning. By lunch, I'm starting to wonder.

So, a key part of UX is denial of service. Keeping them out until they meet the rules for being let in is core to pretty much all modern software. If you are one of those people who reacted to the word 'game' by saying, "Oh, I don't waste my time on games," then start investing some time playing. Angry Birds didn't get big based on the depth of concept or the resulting real-world rewards of mastery. It got big because it tapped into something that made people willing and eager to put themselves through the experience. If we are going to build compelling software, we need to learn how to borrow just enough from the game mechanics to help us along.

Color My World

My wife and I once stood in a computer store in the very early days of the IBM PC and listened to a salesman justify why a green-screen PC was superior to a full-color Commodore 64. "Color is for children. This is a business machine." Yes, there are still, in 2017, plenty of green-screens in business. But nobody is arguing that color has no role in the workplace. Graphs and charts and documentation all work better in color. The web is filled with color. As go blue, red, and all the rest

We need to look at UX with fresh eyes. We need to understand that UX is in every part of what we do. It is not, and should not be, a coat of paint slapped on later.

of the rainbow, everything else that attracts will find its way into some aspect of our business world.

We need to look at UX with fresh eyes. We need to understand that UX is in every part of what we do. It is not, and should not be, a coat of paint slapped on later.

You Knew This Already

Part of UX is naming. When I call a program `INV.6.3`, it might be handy for finding it on the Inventory menu, option six, then option three, but it doesn't communicate purpose. Names are important. Except when they aren't. So, here's a thing you probably know but have never thought about: The web has a huge naming error baked into its core: `index.html`.

It is, other than a few edge cases, what we name the landing page of each part of our site. Granted, depending on our choice of programming language, it might be `index.aspx` or `index.php`, or some other variant, but the key word, the key *name* is `index`. Why? Because the web is HTML and HTML is a subset of SGML and SGML was designed for managing documents and therefore — deep breath, run-on sentences are tiring — what better place to start than a table of contents or, wait for it, an *index*! So, millions of programmers call their landing page `index.<something>` even when that is wildly inaccurate. Assumptions

form design decisions. Design forms a large part of UX.

Apple? Watch!

I have rarely worn a watch since the day I realized I was carrying three other devices that also told the time. I'm a roll-up-the-sleeves guy. Watches get knocked into walls and doors when I wear them. Why put something fragile, expensive, or both, on my flying wrists? I'm sure I'm not the only one that has decided I don't need a face (analog or digital) on my arm.

So, why do the Fitbit, the Apple Watch, and the various bits of Android arm candy exist? Don't we all carry phones with the time on them? Sometimes a laptop or tablet or second phone as well? Hardware is part of UX. For some people having a device that they can turn toward their face without using their fingers is important. The ability to pair it to the phone makes the watch an extension — a feature — of the phone. That's transformative. It isn't another device which tells time, its another interface into the same device. It has become valuable, not for what it does, but for how it lets you interact with what you already use. The watch literally jumped categories and became more relevant in the process.

Beyond IT

UX isn't just about our part of the business world. Imagine trying to get people to routinely spend as much (or more) money on a train ride than on a plane. The train takes longer. It isn't much wider or more get-up-and-walk friendly. How did Amtrak make that work? Well, some of it is due to people who don't like the experience of flying. Some of it came from adding power plugs by every seat in the premium trains. Some of it came from refresh-

ing the interiors of those trains. Adding amenities. All of those things are UX things.

The train experience is different than the plane. The train UX is different. Those are the same sentence. So, remember, the cubicle that you hate and the private office that you love are part of the UX of your job. It isn't just the software on your desk. The mouse (I hate track pads!) or track pad (the mouse hurts my shoulder!) are part of your UX. Maybe, you like the cubicle because you just have to lean back to confer face-to-face with a colleague instead of having to plan a meeting or walk to the other side of the building. All the things you experience come together to form your view of things.

Reign IT In


“That’s great, Chuck, but I just need to make a web page that won’t cause

the user to cringe,” you say. How do we pull back from the totality and focus on our project’s UX? The key words here are empathy and common ground.

If the user only speaks German, maybe your site (or app or software suite) should be displaying prompts in German and help in German and FAQs in German. Find out before you start building it in English or Portuguese. If the users call a job a project, then use the term project. Don’t force them to use your idioms and slang. If half your target base is shade-blind, don’t use tone-on-tone. The key idea here is that until you experience — there’s that word again — how your users see the world, you can only get it right by accident.

Next article we’ll address a topic that has been chewing up a lot of my time

recently: ADA compliance in web development. In the U.S. we have a legal requirement to make accessibility a core part of how we design software. Other countries have similar measures. Oh, and if you work with the U.S. government or work for them, 503 rules are also in effect. Compliance, as Susan Joslyn can tell you, can be fun. **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, Pax Solaria, PerhelionSF, and the Interrogative series, which begins with Tiago and the Masterless.

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