$7.00 U.S.

INTERNATIONAL

Spectrum®

THE MULTIVALUE TECHNOLOGY MAGAZINE | NOVEMBER/DECEMBER 2015

Backups Without The Wait

ALSO IN THIS ISSUE:

- OAUTH 2 Login with MultiValue BASIC
- 3.3 Million Identities: How MultiValue Can Protect Personal Data
- MultiValue Industry End-of-Year Recap
- Business Tech: Ostrich Edition

The Premier MultiValue User Event

INTERNATIONAL Spectrum®
2016 MultiValue Conference

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

RETURNING TO THE WIGWAM , LITCHFIELD PARK, ARIZONA

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

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

Topics covered at the International Spectrum MultiValue Conference include:

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

...and more.

REGISTER AT
www.intl-spectrum.com/conference

INTERNATIONAL Spectrum®
THE MULTIVALUE TECHNOLOGY MAGAZINE

FEATURES | NOVEMBER/DECEMBER 2015

DEPARTMENTS

From the Inside

The exploration and creation of Next Generation Business Software is on the minds of a lot of business software developers lately. Almost all business software is showing its age. Many applications have been around for 15-20 years, but most of our MultiValue applications have been around for over 30-40. Next Generation Business Software is not just about the client. Nor is it just about the user.

I've been watching what a lot of companies are describing as Next Generation Business Software to see what the trends are while watching how people are tackling the problem of the Monolithic ERP, MRP, and LOB software package. For those who aren't familiar with the acronyms:

- ERP: Enterprise Resource Planning.
- MRP: Manufacture Resource Planning.
- LOB: Line of Business, *i.e. accounting, purchasing, invoices systems*.

For example, let's look at one of the most popular small business and personal business applications on the market: Quickbooks. It has been available since 1994. It's old enough to drink in most countries. SAP has been around since 1973 but is likely using a version of a core that was written in 2001. That's 15 years old.

Over the years, there have been a lot of changes to the IT industry. Products and features have been bolted-on to keep up with changes in software and user requirements. But very little was changed in the central code of the applications themselves. Keep in mind, I'm not talking about just MultiValue software, but all business software.

Instead of incorporating each feature into an ERP or LOB application suite, software companies bolted-on the cool new abilities. As a result, the core wasn't enhanced, it was burdened with the new requirements. The core application became more complex with each addition. This complexity causes the applications to become more Monolithic, which also meant it became harder to bolt-on each successive change. To truly create new software, there would have to be a fundamental change.

Next Generation Business Software isn't going to be just about the UI or features or UX or cloud or mobile or Internet of Things, or the hundreds of other things that companies say are important. It's going to be about all of it and none of it.

Any one of these things would require flexibility in the core software of the application; the combination of these things will likely push any well-established software past its limits. Developers are running into a wall when trying to update software because the core is too fragile. A true Next Generation application will not only have functional and flexible user experience, but will be flexible in data and application interfaces in its core as well.

Many existing applications already provide easy access to the data through reports, exports, or data mining. However, they don't tend to provide easy access, updates, and manipulation of the core business data and logic. The data flow

INTERNATIONAL Spectrum
NOVEMBER/DECEMBER 2015

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

http://www.intl-spectrum.com

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

International Spectrum is published six (6) times per year at the subscription price of $40.00 U.S. in the U.S.A.; $45.00 U.S. in Canada and Mexico; $50.00 U.S. for other countries. Single copy rates are $7.00 U.S. in the U.S.A. and Canada, and $9.00 U.S. in all other countries. International Spectrum is published by International Specturm, Inc., 3691 E. 102nd Ct., Thornton, CO 80229; Tel: 720/259-1356; Fax: 603/250-0664 E-Mail: request@intl-spectrum.com.

PRINTED IN USA

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

twitter.com/intlspectrum   intl-spectrum.com/facebook is, relatively speaking, simplistic. It must go in one way and come out the other. All the alternative accesses being built for bolt-ons are adding complexity, and in some cases are adversely restrictive in how information can be manipulated or accessed.

How will this affect MultiValue business software? In reality, we have more flexibility in our core systems than most other business software does. Additionally, we do not have a separation of core business logic from the data, and the data from the core business logic. This provides us with the basis for the kind of flexibility which doesn't exist in other systems, at least not without major rewrites.

This flexibility provides easy access to statically-stored data and derived data — data generated from the business rules applied to the literal data. For example, a price in your inventory isn't just about cost and margins stored in the database, but also about the supporting information like turns, vendor costing, promotional pricing, and ROI. All this is calculated/analytic data. MultiValue allows us to determine this metadata on-the-fly.

Another example of this is order creation. You need to have access to workflows that drive that order through the business systems, not just "how to save the information to disk" so that you can run meaningful reports against it.

Workflows change the data depending on where you are in the process: The creation of an order drives warehousing and inventory data alterations; customer requests or complaints drive inventory controls and clearances; back-orders drive website e-commerce data; e-commerce data and social media feeds drive purchasing of products and future inventory sell-off; raw material purchasing drives product availability and manufacturing production hours; sales data from mobile interfaces drive production rescheduling.

Creating new UI for consumers or employees cannot be done properly if the core application doesn't provide the underlying flexibility to the client interface. This isn't about simply updating data anymore, but about responding to changes in the dynamics which drive the business, and therefore, drive the application logic.

In this article, I've hit on a lot of little topics that might seem separate, but in reality are linked together to make a game-changing whole. I will be talking about the overall concept in more detail during the upcoming 2016 Conference, so make sure to join me there.

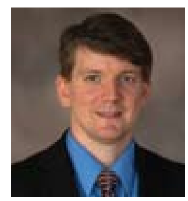

NATHAN RECTOR
President
International Spectrum
nathan@intl-spectrum.com

Wait! Backup!

BY KEVIN KING

I hate to wait. Standing in lines at the supermarket, or heaven forbid the Department of Motor Vehicles, is a pain I'd rather avoid. But even more than the DMV, I hate to wait for a paused Unidata system to come back online when the backup is running. I *can* avoid the DMV, except once every few years. With Unidata backups, however, that delay that can happen daily. Even more irritating, it may be entirely unnecessary!

Pausing your database to take a quiescent copy is a good thing. Historically, backups have paused the database for the entire time that the backup is running. With the newer systems that's no longer necessary. For example, one can pause/lock a Unidata/Universe database environment, then take a snapshot or split a mirror, then resume/unlock the database, and then run the backup off the — secondary — quiescent copy. This can reduce your total downtime from several hours to a couple of seconds or less.

For those using AIX, IBM has this backup strategy known as "split backup off mirror" (SBOM). With this backup strategy, the file system is mirrored using the logical volume manager. The backup procedure pauses the database and then splits off a mirror from the volume group, mounting it separately as a copy of the original file system. Once this mirror has been split, the database on the original drives is resumed and the backup commences from the copy. When the copy is complete, the mirror, previously split, is re-added to the volume group and re-synced to the other copies. The database is paused for merely a couple of seconds, rather than the hours it would take for the full backup procedure.

A snapshot backup is really quite easy to do, assuming the logical volume manager has been configured with a little bit of unclaimed storage. Here's the 40,000-foot view of the process...

A similar thing can be done with Red Hat Enterprise Linux (RHEL) as well, although it's not the most efficient option. Instead, on RHEL a snapshot can be taken of the file system, mounted, backed-up, and then when the backup is complete, the snapshot can be removed. There are a couple of benefits to this approach:

- Reduced storage: When a snapshot is taken on RHEL, the snapshot only needs enough storage to hold any changes to the original filesystem to bring the snapshot up-to-date with the original copy. If the snapshot is done at a time when the system is not under load, the amount of disk needed for the snapshot can be very small. By contrast, a mirror will persistently need 100% of the original drive space.
- No re-sync on the snapshot: When the snapshot backup is complete, the snapshot can be removed. It is merely a frozen moment in time, so by removing it you only lose that moment. Everything that happened since the snapshot remains intact on the original copy. And removing a snapshot is remarkably fast! With the SBOM approach, it can take hours for the mirror to be joined and updated with the original copy. The system isn't down – these are

not hours of downtime, per se – but they are still hours of I/O that can impact the performance of the rest of the system.

A snapshot backup is really quite easy to do, assuming the logical volume manager has been configured with a little bit of unclaimed storage. Here's the 40,000-foot view of the process:

- Pause the database
- Take a snapshot of the file system
- Resume the database
- Mount the snapshot somewhere in the file system
- Do the backup
- Remove the snapshot

It's also important to add in some basic validation to ensure that you don't start a backup when another backup is running; that the database isn't already paused, that sort of thing.

Here are the commands to set up a snapshot backup on UniVerse and UniData systems:

Step 1: Pause the Database

On Unidata, the `dbpause` command run by someone with superuser status will pause the database. This allows programs to read information from the files, but any attempts to write anything will block (hang) until the database is resumed.

On Universe, the comparable command is `uv -admin -L` to lock the database from being updated.

Step 2: Take a Snapshot of the Database

The `lvcreate` command is used to create the snapshot using this syntax:

```
lvcreate -LsizeG -s -n snap /dev/
volumeGroup/volumeName
```

The *size* parameter defines how big the snapshot volume needs to be. Typically, 10% to 20% of the size of the original drive will suffice. However, that is merely a SWAG [*educated guess - Ed*] and should be adjusted based on the updates happening while the backup is running. It's important that the snapshot is big enough. Otherwise, during the backup, the snapshot could become invalidated by updates. However, it doesn't need to be excessively large.

The *volumeGroup* parameter defines the volume group where *volumeName* can be found. The `lvs` command can be useful for reporting the logical volumes currently defined. If, for example, your volume group is called newsystem and the logical volume is called root and it's 200 GB in size, a snapshot sized at 10% can be created with:

```
lvcreate -L20G -s -n snap /dev/
newsystem/root
```

Note that the newsystem volume group will need to have 20 GB in unclaimed storage to allow this snapshot to be created. The `pvs` command can be used to see how much free space exists on the physical volumes used in the logical volume group.

Step 3: Resume the Database

On Unidata, the database can be resumed by a superuser using the `dbresume` command. On Universe, this command can be used in the operating system to unlock the database:

```
uv -admin -U
```

Step 4: Mount the Snapshot Somewhere

This is where things can get a little tricky. The default filesystem on RHEL is called xfs. Each xfs storage volume contains a unique ID. Because the snapshot is a moment in time of the original volume, both the snapshot and the original copy have the same (should be) unique ID. We can mount the snapshot by telling RHEL to ignore the duplicate ID, but we do that to our peril. If the system goes down when the two volumes with the same unique ID are mounted, there's a strong likelihood that the system will not come back up again without some assistance.

To mitigate this, we'll mount the drive temporarily so that some final updates can be done, and then unmount it right away. There is still a risk that the system might go down in between these two operations, but that risk is minimal as both commands will finish nearly instantaneously. Then we'll tell the xfs administrator to generate a new unique ID for that drive. And finally, when that's done, we can mount the drive somewhere that the backup process can see it. The following commands do this:

```
mount -n -o nouuid /dev/
volumeGroup/snap /mnt/snap
umount /mnt/snap
xfs_admin -U generate /dev/
volumeGroup/snap
mount /dev/volumeGroup/snap /mnt/
snap
```

In this scenario, I'm mounting the snapshot volume at `/mnt/snap`. You can, of course, mount the snapshot anywhere you want, as long as the directory exists prior to the mount command.

In the first `mount` command, the `-n` option says this is a temporary mount and the `/etc/mtab` file should not be updated by the mount. This is one attempt to mitigate the issue where there are two volumes mounted with the same unique ID. However, I found that even with the `-n`, the `/etc/mtab` file is updated so this option may not be as helpful as intended.

The `-o nouuid` option on the first mount command says to ignore the fact that the two volumes have the same ID, and mount it anyway. During the mount, the logical volume manager does some cleanup on the mounted volume to ensure that the drive will later mount cleanly. Because this is the riskiest command of the bunch, with the potential to leave two volumes mounted with the same ID, it is immediately followed by a umount to unmount the snapshot volume.

The `xfs_admin` command with the `-U generate` option will generate a new unique ID for the snapshot volume. Once this is done we can then mount the snapshot drive again — cleanly, safely — and without requiring any options.

Total elapsed time for these four commands is a second or less, and they generate so little overhead on the system that normal everyday users won't even know all this magic is happening.

Step 5: Do the Backup

Once the backup drive is mounted, you can use whatever tools you have to do the backup. Maybe you're backing up to an appliance ala Unitrends or Synology, maybe you're `rsync`ing the data to off-site storage, or maybe you're backing up to tape. It really doesn't matter how you backup, except to say that you should be backing up to something external to the current system and preferably to something off-site, like the cloud storage replication option available on Unitrends appliances.

Step 6: Remove the Snapshot

Once the backup is complete, there's no point to holding on to the snapshot. The `lvremove` command can quickly dispatch it, freeing that storage space. However, before removing the snapshot logical volume, it must first be unmounted:

```
umount /mnt/snap
lvremove -f /dev/volumeGroup/snap
```

Of course, as in all previous commands, replace *volumeGroup* with the appropriate volume group for your system.

If you're thinking it might be good to hold on to that snapshot "just in case," allow me to talk you out of that idea. Remember, we only sized the snapshot volume to support the changes that would be happening to the original during the backup. The longer the snapshot sticks around, the more likely it is that you'll overflow that space and invalidate the snapshot. More importantly, when there's an active volume snapshot, everything that happens to the original volume gets logged for the snapshot, resulting in double updates. Deleting the snapshot frees up that storage and allows the system to go back to single updates for the original volume only.

This is notably a high-level view of the process to do a snapshot backup, but hopefully it illustrates that this is not all that difficult of a procedure. Wrap these commands up in some validation to ensure that nobody does anything stupid, and watch your backup downtime decrease to virtually nothing! IS


KEVIN KING is the President and Chief Technologist with Precision Solutions, Inc., a leading technology solutions provider in Longmont, Colorado. He can be reached by email at Kevin@PrecisOnline.com or by voice at 303/651-7050.

OAUTH 2 Login with MultiValue BASIC Part 2

BY NATHAN RECTOR

Once you have your project created, and you have signed your life away via Google's "Terms and Agreements" dialog, you can get your keys from the "Credentials" sub-menu under the "APIs & Auth" section by adding your credential request.

In the last article, I gave an overview of OAUTH 2 and the different ways login can be done. In this article I'll talk about some of the data files and routines you need to setup to make this all work.

The Control Record

Part of what makes OAUTH 2 work is the additional information that is sent along with each URL request. Each request requires a token to be included but – to make things more interesting – your system will also need to keep track of other security information in order to request the current tokens. This data needs to be available for OAUTH 2 to work correctly.

On my systems at International Spectrum, I have created a control file call OAUTH.CONTROL.SUPT. The control record will contain information needed when requesting new Auth Tokens to be sent to the web service. It also includes the rules for accessing these web services. The record structure is found in below.

```
001 - Consumer Key
002 - Consumer Secret
003 - Token Request URL
004 - Token Access URL
006 - Final Auth Token
008 - Refresh Token Returned
009 - Session Handle Returned
010 - Date Authorized
011 - Time Authorized
012 - Token Expires In (Seconds)
013 - OAUTH Version (2.0)
015 - Request Token (Internal Uses)
016 - Request Secret (Internal Uses)
017 - Request Pin Number
018 - Request Pin URL
019 - User Name
020 - User Password
021 - OAUTH Login Flow
022 - Extra OAUTH Field Names (Returned)
023 - Extra OAUTH Field Values (Returned)
025 - Redirect URL
```

Much of the information in this control record is a combination of transactional, or transitory, data only needed during the authentication process, and fixed parts that should always be unchanged. Since this file contains passwords and identifying information, you may decide to encrypt the individual fields or just encrypt the whole file using "Encryption at Rest."

Another thing to keep in mind is that your may have more than one OAUTH 2 control record per web service. A classic example of this is the Salesforce REST service. If you need your applications to update to Salesforce separately for each user, then you need to have a control record for each individual user. This may be required for audits and compliance.

You can see an example of what information is needed, and how Google+'s OAUTH 2 uses it in below.

```
001 xxxxxxxxxsbsqnupkl3p0bhrievchdo8j58.apps.
    googleusercontent.com
002 tGkk2QnhKbtpxxxxxxxx
```

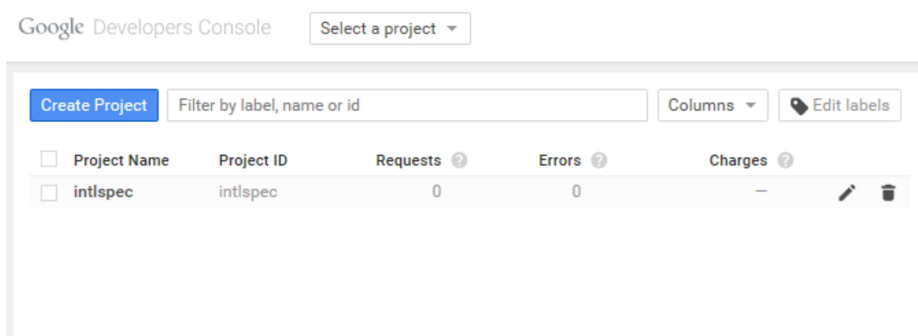

Fig. 1

```
003 https://accounts.google.com/o/OAUTH 2/
    auth?scope=email profile
004 https://accounts.google.com/o/OAUTH 2/token
013 2.0
021 native
```

Where Is All This Stuff?

Before you can do anything, you have to sign up with the web service you wish to authenticate with in order to get your consumer key and secret. Since I'm using Google here, and you'll likely want to try these examples as well, let's do a quick run through to get access to the Google API.

To get started with Google APIs, you need to go to Google's Development Console (*Figure 1*): https://console.developers.google.com/project

Google has an in-depth help document that will take you through all the steps at https://developers.google.com/+/web/api/rest/oauth#acquiring-and-using-an-api-key

Once you have your project created and you have signed your life away via Google's "Terms and Agreements" dialog, you can get your keys from the "Credentials" sub-menu under the "APIs & Auth" section by adding your credential request. Be sure you are adding an "OAUTH 2.0 Client IDs" request. There are other kinds, but they are not needed for this example.

Now comes the bit where we code.

The HttpClient

Having done the manual steps, we now need to teach our computer how to talk with the web service. HttpClient is the communication tool for this job. Just to be clear, I am not talking about a web browser. They are not the same thing.

A web browser is a program that is used to render HTML. The browser does have an HttpClient inside of it, but just as an engine isn't a car, the two are not synonymous. HttpClient is a connector — the middleware — that links to a web service, asks for something, retrieves it, and then closes the connection. The information needs to be passed off to another program — for example: a web browser — which can do *something* with it.

There are lots of different options for building an HttpClient. It can be created using raw socket calls or using a Telnet client. All forms of MultiValue databases have a way to either establish an HttpClient or provide an interface with the tools needed to create an HttpClient.

You can find examples of source code for building HttpClients at:

http://www.intl-spectrum.com/resource/category/190/httpclient.aspx.

One tried and true solution is to use a command-line tool called cURL. A version of this command-line tool can be found on Windows, Linux, AIX, and most any other O/S. If you visit the URL above, you'll see an article and source code for using cURL as the HttpClient.

Once you have your HttpClient, you will need to wrap it into a new program. This code example can be found at the above URL, and is called IS.HTTPCLIENT. For the purposes of this article, I'm going to create a new program called IS.OAUTH.HTTPCLIENT with the same parameters as the example, except that I'm going to add an additional one called OAUTH.ID, which will contain the ID of the OAUTH.CONTROL.SUPT record.

SUBROUTINE

```
IS.OAUTH.HTTPCLIENT(OAUTH.ID,HTTP.METHOD,URL,HEADER.
DATA,POST.DATA,RESP.HTTP.STATUS,RESP.HEADER.DATA,RESP.
DATA,HTTP.ERROR)
```

The first thing that the program will need to do after reading the OAUTH.CONTROL record, is to check to see if there is any Bearer information to be added to the HTTP request. The Bearer information is the OAuth Token returned from the OAuth Login service which is stored in OAUTH.CONTROL<6>.

If you have Bearer information, you will need to check to see if it has expired using the Auth Date (in OAUTH.CONTROL<10>), Auth Time (in OAUTH.CONTROL<11>), and ExpiresIn Seconds (in OAUTH.CONTROL<12>). See an example of how to do this below.

```
HAS.EXPIRED  = 0
*
BEGIN CASE
 CASE (ITEM.OAUTH<CONTROL$ACCESS.TOKEN> EQ ‘’)
 HAS.EXPIRED = 3
 CASE (ITEM.OAUTH<CONTROL$DATE.AUTH> EQ ‘’)
 HAS.EXPIRED = 2
 CASE (ITEM.OAUTH<CONTROL$TOKEN.EXPIRE.IN> LE 0)
 HAS.EXPIRED = 0 ;* Never expires
 CASE 1
 AUTH.EXPIRE   = (ITEM.OAUTH<CONTROL$DATE.AUTH> *
86400)
 AUTH.EXPIRE  += ITEM.OAUTH<CONTROL$TIME.AUTH>
 AUTH.EXPIRE  += ITEM.OAUTH<CONTROL$TOKEN.EXPIRE.IN>
 *
 EXPIRE.CHECK  = (DATE() * 86400)
 EXPIRE.CHECK += TIME()
 *
 IF (EXPIRE.CHECK > AUTH.EXPIRE) THEN
 HAS.EXPIRED = 1
 END
END CASE
```

If the Token has not expired, then you need to add it to the HTTP headers so that IS.HTTPCLIENT can pass it along

```
LAST.HEADER.POS = DCOUNT(HEADER.DATA<1>,@VM) + 1
HEADER.DATA<1,LAST.HEADER.POS> = ‘Authorization’
HEADER.DATA<2,LAST.HEADER.POS> = ITEM.
OAUTH<CONTROL$ACCESS.TOKEN>
*
CALL IS.HTTPCLIENT(HTTP.METHOD,URL,HEADER.DATA,POST.
DATA,RESP.HTTP.STATUS,RESP.HEADER.DATA,RESP.DATA,HTTP.
ERROR)
```

Next...

Now that you have the basic structure and requirements to handle an OAuth web service, the next article will talk about the handshaking needed to acquire the Bearer token. IS

International Spectrum 2015 End-of-Year Recap

Entrinsik

Q: ***What new features or services did Entrinsik introduce in 2015?***

2015 was a busy year in development at Entrinsik as we introduced versions 4.5 and 4.6 of our Informer BI & reporting software. Enhancements included customized formatted output, embedded links, new maps for dashboards, simultaneous reverse link creation, and much more. MultiValue users can now get even more out of their Informer system with these new reporting, data analysis capabilities, and interactive dashboard features. Informer Roadshows also traveled around the country throughout 2015 offering advanced training for Informer users, ensuring that customers are leveraging all the benefits of the software.

Our predecessor software product, Semtek, went through a significant change/upgrade, becoming Enrole SaaS. The MultiValue database provides many solutions for our customers and corporate clients whose business needs are not simple or generic, and coupled with an intuitive and clickable dashboard user-interface it became an attractive and exceedingly easy to use total solution platform. Training reduced from five or more days in the past down to two, thus proving that the time taken to rewrite and redevelop the entire product and processes was well worth the substantial effort and investment.

Q: ***What was the greatest or most exciting thing that happened to your company or product(s) in 2015?***

The most exciting thing to happen to Entrinsik in 2015 was its inclusion in some important benchmark reports and award contests. Entrinsik won the SugarCON 2015's *Best in Class* award for Data, Information, and Analytics in April of this year. Entrinsik Informer was also included both in Gleanster's data visualization benchmark report with a ***Best*** ranking in all four categories, as well as in the BARC BI Survey 15 with ***Leader*** rankings in six categories. Additionally, Entrinsik was honored to be included in DBTA's 100 *Companies that Matter Most in Data* for the third consecutive year.

It was an exciting year for our Enrole SaaS product as well! The beginning of 2015 saw Enrole's first customers live and running on a fully cloud-based, web application, for the non-credit departments of both Professional and Continuing Education entities. We gained seven orders in eleven months with more in the pipeline for early 2016. Providing outstanding customer service and attention to detail for our customers has given Enrole substantial credibility in the education arena, but also challenged us with innovative new procedures, all of which have successfully been incorporated into our product and policies.

In addition, both Entrinsik products, Enrole and Informer, were rated highly by users of our software on sites like IT Central, TrustRadius, and Capterra, independent sites that provide user-generated reviews of business software.

Q: *What would you like Spectrum readers to watch for in 2016?*

Entrinsik has some very big and exciting things coming in 2016. Readers should be watching for updates and changes in Informer's interface, architecture, features and capabilities as we are continuously evolving Informer to be better than ever before.

Current Informer development efforts are focusing on upgrades to the user interface and architecture, plugin usability, and overall ease of use for non-technical users. Enrole development is ongoing and focused on some features from the predecessor software, and applying new functionality and integrations, identified by existing customers and new prospects alike. Another busy year ahead is planned!

Kore Technologies

Q: *What new features or services did Kore Technologies introduce in 2015?*

We announced KommerceServer eCommerce Suite Release 5.4. It included many performance improvements and new features. The search and catalog browsing engine was significantly improved to deliver faster results and be more scalable and consistent. The Product Quick Search feature was enhanced to deliver faster and more relevant results to search inquires. A new Request for Quote (RFQ) capability enables sales teams to engage with clients without disclosing pricing and can integrate to the back-office ERP or to Client Relationship Management (CRM) applications. A Branch and Agency locater capability was added enabling the implementation of loyalty programs by generating Web traffic for key customers and/or affiliates. A Product Inventory Search by Branch helps customers find inventory availability at specific locations. A new Coupon and Promotion Code feature helps market your business to new and existing customers. Freight Carrier Zones and Rates Management improvements make it easier to manage internal truck fleets as well as other shipping carriers.

Kore's Kourier Integrator data warehousing and enterprise integration solution was used to integrate with many popular on-premise and cloud-based applications including: Bamboo HR, Blackboard, OnBase and ServiceMax. These new integrations show the power and breadth of capabilities available with Kore's flagship product.

Q: *What was the greatest or most exciting thing that happened to your company or product in 2015?*

Our Kourier Integrator data warehousing and integration product continues to make significant inroads into the distribution and higher education markets. Kourier's rapid deployment, high return on investment and expansive feature set for extract, transform, load and enterprise application integration (such as the ability to aggregate data from multiple disparate data sources and real-time integration via RESTful Web Services) make for an industrial strength solution that has become well accepted in the marketplace.

All the product improvements made in KommerceServer this past year helped Kore make significant inroads into the distribution market with additional implementations of our eCommerce solution for B2B and B2C distributors. A number of clients are rolling out new websites with the enhanced capabilities of KommerceServer.

And of course, success leads to more success. It is very exciting for Kore Technologies to have one of its products selected as a Trend-Setting product by Database Trends and Applications magazine! This is the third year in a row for Kore, and the second year in a row that KommerceServer has earned this honor.

Q: *What would you like Spectrum readers to watch for in 2016?*

Kore plans for the next release of Kourier Integrator to include rolling out more real-time integration capabilities with significant enhancements to our RESTful Web Services. Expanded support will be provided for non-MultiValue data sources next year with the addition of new data sources, a new Quick Start Workbench for non-MultiValue databases, and new data input/output options (such as JSON).

KommerceServer eCommerce Suite will be providing improved support for mobile sales teams, enabling them to more easily do business on the road. Our mobile strategy for next year will include an improved responsive web design and optimized architecture for mobile. Our eCommerce clients can also expect an improved search engine that will offer higher performance and advanced searching capabilities.

Ladybridge Systems

Q: *What new features or services did you introduce in 2015?*

This year has seen a large number of minor enhancements as requested by

QM users and a few more significant additions.

The scope of the connection pooling feature referenced as a 2015 development in last year's recap article has been extended to allow its use with both QMClient sessions *and* phantom processes. Use of this feature can give significantly improved responsiveness for short life processes such as web servers.

The multi-tenanting features that allow multiple totally isolated instances of QM on a single server have been enhanced to offer improved security capabilities.

The data encryption features of QM have been extended to allow encryption of data in directory files and indices.

The final release of 2015 includes the first phase of support for simple web services.

Q: ***What was the greatest or most exciting thing that happened to your company or product in 2015?***

In January 2015, Ladybridge Systems announced the appointment of Zumasys as our worldwide sales and marketing partner for OpenQM. This relationship has, as expected, brought more visibility to the product, greatly increased the resources available to handle customer inquiries, and seen some important new sites come online. Previously established QM resellers continue to operate largely unchanged.

Q: ***What would you like Spectrum readers to watch for in 2016?***

QM continues to develop, not only by adding new interfaces but also by extension of the core programming environment. The live backup system that allows creation of a consistent backup whilst the server continues to process new work is scheduled to be released early in 2016. There are many smaller but important developments in the pipeline too.

OnSystex

Q: ***What new features or services did Onsystex introduce in 2015?***

This year we introduced multiple new products and capabilities to enable a broader range of companies to leverage their MV investments while introducing critical improvements in user experience, service orientation, cloud integration and reporting.

OASYS|OMNI Integration Framework is an MV-based architecture which exposes MultiValue business logic, referential integrity, and data as web-based REST APIs. It provides a web framework to design and implement services, expose services to external applications, call the services from existing MV code and delivery power web/mobile interfaces to the MV application.

OASYS|BRIDGE Data Warehouse is a configurable near-real-time SQL replication engine. The platform leverages our OASYS|SERVER technology and provides data discovery, mapping, schema generation, triggered replication, high-performance ingestion services and robust error and service level management capabilities. The platform also contains options for event based replication, net change replication and the ability to configure updates from the SQL data in support of new application development.

OASYS|WEB brings together a number of products, including OASYS|OMNI to enrich green screen applications with web-based capabilities. OASYS|Web|Emulator delivers a traditional 80x24 experience in a web-browser and is designed to lower the cost and broaden the users that can gain access the MV application via the web. OASYS|Web|Forms is advanced screen-scrape technology which converts 80x24 emulations into .NET code which in turn produces a responsive web-design experience than can be modified without changing host screens.

Q: ***What was the greatest or most exciting thing that happened to your company or product in 2015?***

Our OASYS|OMNI launch was quickly adopted and put into production in over 20 sites in its first year and we expect that number to grow to 100 in 2016. One MV developer now using OASYS|OMNI decommissioned dozens of .NET and other scripted API's said, "OMNI is a game-changer for opening up our UniVerse system."

The introduction of our OASYS|WEB solutions allows MV applications to run seamlessly in a web browser including all mobile devices. The ability to split the MV user interface from MV application logic is a game-changer in our industry.

OASYS|BRIDGE a hybrid of OASYS|SERVER and has allowed our customers to build a bridge between MV and SQL databases, the surfacing of Multi-Valued data as normalized and openly accessible remains a core competency.

Q: ***What would you like Spectrum readers to watch for in 2016?***

Using OASYS|OMNI we will release a real-time ODATA integration layer in Q1 which dynamically delivers full Create, Read, Update, and Delete operations understood by many technologies. It will include mapping capabilities to provide specialized views of data across multiple files and business processes.

Additionally, we will be providing other data source implementations optimized for use in real-time business intelligence solutions such as PowerBI and Tableau.

With many of our clients asking about the cloud, we are planning a number of other product extensions designed to easily begin migrating shared application services to the cloud while maintaining tight integration with existing system.

Pick Cloud

Q: *What new features or services did Pick Cloud, Inc. introduce in 2015?*

This year the Zayo Group (NYSE: ZAYO) purchased our data center partner, Latisys. This enabled us to not only to leverage the rise in the number of data centers our customers now have access to across the US and Europe but also to increase our offerings from Windows and Linux to include other operating systems as well.

Our partnership with Zayo, as well as Rocket Software, Ladybridge (OpenQM DBaaS) and SDSI (CirrusPrint) has allowed us to be the front-runner in offering the most cost-effective cloud offering in the MultiValue market.

While we continue to offer our flagship products such as OpenQM DBaaS and short-term server leasing (both starting at $15!) and cloud printing products like CirrusPrint, we are always searching for new products and partnerships that can enable us to make cloud computing available to everyone – even the small to medium businesses (SMBs) who felt the cost was out of their reach.

Q: *What was the greatest or most exciting thing that happened to your company or product in 2015?*

This year marked the 50th anniversary of PICK/MultiValue. Its longevity is a true testament to the technology my father helped to create. We participated in many MultiValue conferences and user groups, from the inaugural rejuvenation of the Florida Pick Users Group, CMUG, and OCMUG, Rocket's MVU, to the International Spectrum Conference. We spread the word about cloud computing as well as the history of PICK and its future.

In 2015 we were again included on DBTA's Top 100 Companies that Matter In Data list. This was the third consecutive year we were included on this list with such great MultiValue companies as Rocket Software and Entrinsik.

Q: *What would you like Spectrum readers to watch for in 2016?*

Pick Cloud, Inc. believes in preserving and rejuvenating MultiValue. This has been our mission statement since our inception. We believe that after 50 years the PICK/MultiValue database has withstood the test of time and through our partnerships and MultiValue expertise, we continue to offer the best, most cost effective, cloud hosting solution for the MV community.

Continue to look for us as we bring more awareness to the MultiValue community of cloud computing and how it's the perfect fit for PICK and MultiValue applications! We will also continue our support of the MultiValue community in the form of PICK MultiValue User Groups and trade conferences.

Pick Cloud, Inc. will continue providing more products and services that are cost effective for MultiValue application owners and companies of any size.

Rocket Software

Q: *What new features or services did Rocket Software introduce in 2015?*

Rocket Mobile for MultiValue: Someone without previous mobile development experience can quickly build new mobile applications — no special coding skills required with Rocket Mobile for MultiValue. Drag-and-drop widgets and templates simplify development. Leverage MV RESTful services to accelerate integration with and reuse for your MultiValue application.

Rocket UniData 8.1 marks the largest release in terms of enhanced features for the application platform in over ten years. With a true 64-bit architecture, 64-bit file support, security enhancements, account based licensing, support for IPv6, install/upgrade improvements, replication performance monitor, and local subroutines/functions with variable scoping.

Rocket CorVu NG is now integrated D3 DBMS. CorVu NG enables development teams to rapidly build and deploy comprehensive data visualizations with real-time data access, interactive dashboards, and drill-down. CorVu NG provides the right information wherever it's needed — to the desktop, web or mobile device and is integrated with U2 and D3 families of MultiValue databases.

Self-Paced Training is available for U2 Fundamentals for Rocket UniData and UniVerse beginners. Because the class also has an eLab environment with it, allowing you to practice and do exercises while learning, there is a need to limit the time-frame in which you take the class. For more information visit: https://www.rocketsoftware.com/u2-self-paced-training-your-schedule.

Q: ***What was the greatest or most exciting thing that happened to your company or product in 2015?***

Rocket MultiValue University 2015 in Las Vegas reinvigorated the passion for MV technology within our community. This international event attracted attendees from five continents and included not only the U2 but also the D3 family of products. Members of our MultiValue community came together to share and learn more about the latest developments and emerging technologies associated with Rocket MV.

Besides Mobile Enablement, MVU attendees also learned more about modern integration of other languages with MultiValue, such as Python, about security, and high availability. The benefits associated with our latest product releases and other Rocket technologies such as Business Intelligence, Secure File Sharing, Application Lifecycle Management, and much more were also discussed.

Q: ***What would you like Spectrum readers to watch for in 2016?***

In 2016, Rocket will continue to lead business applications into the future of MultiValue by continuing to develop technologies integrated with MultiValue. Python integration with both UniVerse and UniData will be released to complement the native Basic language. Audit Logging will be offered for U2 (first UniVerse and then UniData) to meet the needs of recent auditing demands and help with compliance regulations. Rocket will also offer Self-Service BI with MultiValue data for business users who need to quickly connect to data, create rich visualizations that expose trends, outliers, and historical data, and derive insights for better business decisions. Rocket will also introduce a brand new product that monitors, detects, and cures the system and OS associated with the MV application platform.

SJ+ Systems Associates

Q: ***What new features or services did SJPlus introduce in 2015?***

Every year brings new risks, new threats, new regulations to the world of software development and data processing. But each of these also bring new opportunities to make our software better. PRC gets smarter so that it can put the tools into your hands to create a sleeker, more automated, and more integrated development environment. This results in better software and feedback for *your* users.

This past year has been a big one for social interaction — PRC has been learning to be more sociable, too. Playing well with others means integrating desktop applications, such as editors and compare tools, automating email, managing movement between servers, and attaching documents and test scripts.

Agile is a buzzword that is simply not going to go away. Some companies have been tweaking their use of PRC to support their new agile practices, which has been fun and rewarding. IS

Data Encryption at Rest

MARK FULLER, NORTHGATE INFORMATION SOLUTIONS

Files containing wireless contracts, customer names, customer addresses, and credit card details were found in a recycling bin. Personal data — *3.3 million names, addresses, dates of birth and Social Security numbers* — have been exposed due to stolen portable media devices.

Do I have your attention now?

Data security issues should be of concern to anyone managing data, particularly when it contains PII (Personally Identifiable Information) or credit card information. This article will show you how Reality solves these concerns and minimizes the risk of data falling into the wrong hands using its Data Encryption at Rest functionality.

> ***Reality solves these concerns and minimizes the risk of data falling into the wrong hands.***

What is Data Encryption at Rest?

This term is used for any data that is physically encrypted and stored on some sort of medium — such as a hard drive, thumb drive or even a magnetic tape — and is decrypted on-the-fly when an authorized user unlocks the data using a defined key to access it.

Reality's Data Encryption at Rest is integrated into its user-based security system, making it very easy to administer. We provide tools for managing the keys that unlock your encrypted data. So let's start by reviewing that security model.

User-Based Security Model

The user-based security model requires each user to have a unique user ID. This ID is linked to a security profile that is typically linked to function or department — such as accounting or payroll — however, a profile could exist for each unique user. Therefore, there may be multiple users with IDs for the same security profile, and a user might be connected to multiple security profiles [*Figure 1*], similar to MS Windows users and groups' functionality. It's within the security profile that the encryption keys are stored.

The security system is all controlled via the TCL command `SSM` which is accessible from the system administrators account — SYSMAN [*Figure 2*].

Create an Encryption Key

Now that we have a basic understanding of the security model, we can begin to discuss the steps to create an encryp-

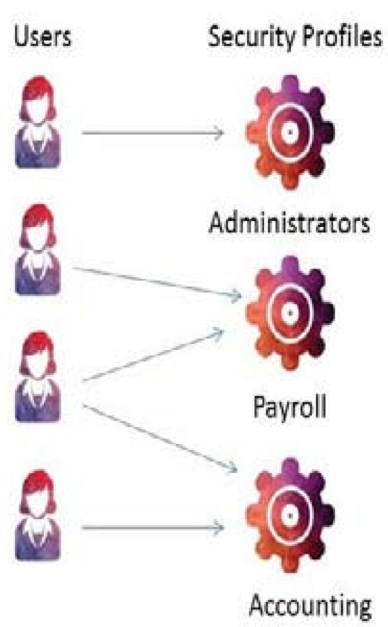

Fig. 1 Reality's Security Model

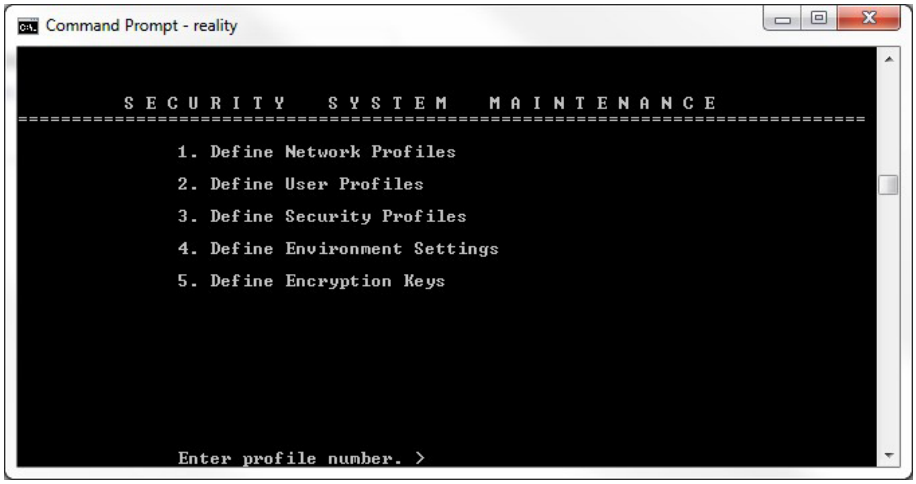

Fig. 2 SSM — Accessible from the SYSMAN Account

tion key, and then how we may use that key to encrypt a file.

The first thing we must do is to set the admin key, which is used for protecting any keys that you define. You do this via the TCL command `SET-ADMIN-KEY  key` where `key` is a sixteen-character encryption key. You will need this key in emergency situations, so it's your responsibility to ensure that the key is stored securely but still accessible to relevant members of staff.

Having identified the files that need to be encrypted, we first have to define the encryption key(s). We do this by using option five of `SSM` [*Figure 2*]. Firstly, we must supply a key ID – this is not the actual key but a reference to it. It is this reference to the key that's used to encrypt or save the file. In *Figure 3* you can see I have created a key called PAYROLL, its type is DES3 (Triple DES) and I have supplied the sixteen-character key with which our data will be encrypted against. Once we file this, for additional security, the key itself is encrypted. Reality supplies an additional tool called `SAVE-REK` that can be used to make a backup of your encryption keys so that you can get them back in an emergency, or use them when transferring encrypted data to another server. Similarly, there is a `LOAD-REK` command to restore your encryptions' keys.

Using an Encryption Key

In order for users to use this key, it must be assigned to their security priority via option three [*Figure 2*] — Define Security Profiles. Within this, field twenty-five — Reality Encryption Keys — must be modified to add our newly created key to it, and we use the key ID and add it to the profile. Remember, there may already be encryption keys

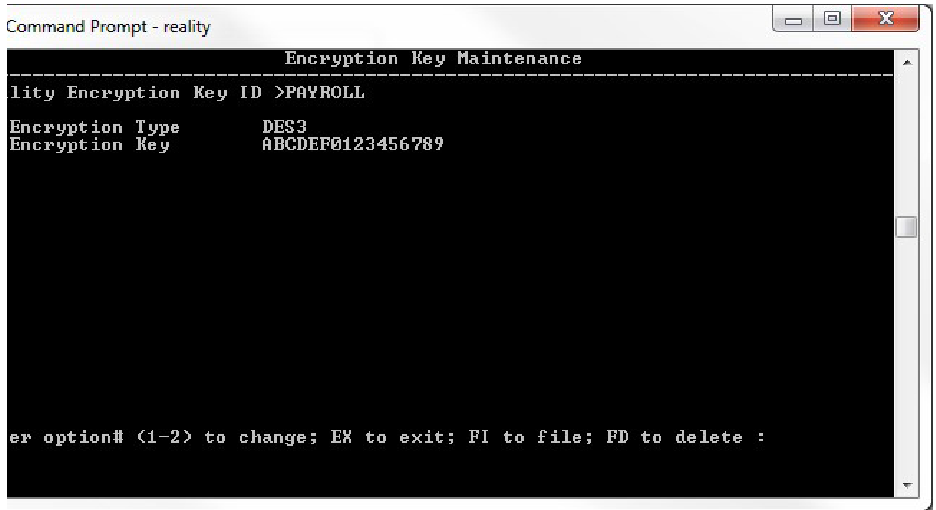

Fig. 3 Defining a New Encryption Key.

in use, as a profile can have multiple keys.

Once you have created your encryption key items, you can create encrypted files to hold your data. This is done by using `CREATE-FILE` with the (C option. You will be prompted for the ID of an encryption key. See the example in *Figure 4*.

```
:CREATE-FILE MYENCFILE (C
Enter Reality encryption key: PAYROLL
[417] File 'MYENCFILE' created.
D code =DL, modulo = 1, separ = 1
[417] File 'MYENCFILE' created.
D code =DL, modulo = 1, separ = 1
```

Fig. 4

Note that by not specifying the modulo and separation from the `CREATE-FILE` we create an auto-sizing file, the modulo and separation is not really `1` and `1`.

```
:CREATE-FILE TEMP (C
Enter Reality encryption key: PAYROLL
[417] File 'TEMP' created.
D code =DL, modulo = 1, separ = 1
[417] File 'TEMP' created.
D code =DL, modulo = 1, separ = 1
:COPY CUSTOMERS *
TO:(TEMP
13 items copied.
:COPY DICT CUSTOMERS *
TO:(DICT TEMP
8 items copied.
:DELETE-FILE CUSTOMERS

:RENAME-FILE TEMP
TO:CUSTOMERS
```

Fig. 5

If you want to encrypt the contents of an existing file, you must create a new encrypted file and copy the contents of the existing file into it (remember to copy the dictionary of the file separately). You can then delete the existing file and rename the new one to have the same name as the deleted file. *Figure 5* shows an exampe to encrypt the data in the file CUSTOMERS:

The file that formerly resided in the `ACCOUNTS` account and was called `TEMP` still resides in the `ACCOUNTS` account but is now called `CUSTOMERS`.

A cautionary note: As with other databases, a simple delete of data does not physically remove the data from the medium. All it does is set a flag to say the data no longer exists. As a result, the steps above will still leave a footprint of the unencrypted data (the old version of the file) on the disk, or other medium, until that space is re-used. So, if you lose the medium, the potential is there to still access unencrypted deleted data.

One solution you can use is to copy the data as above to your newly created encrypted file and, before deleting the original unencrypted file, physically rewrite all the data in the file with a random set of data. You must ensure that the item sizes remain the same. Do not be tempted to write binary data over the existing records, as this will force the item to become 'out of group' again, potentially leaving the original footprint of the item in place. Once that overwrite has been completed you can then delete the file. There are various techniques and methods available to make reading overwritten media harder to get back, and you should be able to find more details on these with a Google search.

Encrypting Individual Fields

If you wish to encrypt certain pieces of information in each item, and control this directly from your application, then you may use the DataBasic function `ENCRYPT`. To decrypt you would use the function `DECRYPT`.

For example:

```
MYKEY = "2":@AM:"0123456789012345"
TEST = ENCRYPT(myfield, MYKEY,5)
```

This would encrypt the data held in the variable `myfield` using Triple DES (represented by the 2) against the key `0123456789012345` and return the data into the variable `TEST`.

Creating an Encrypted Backup

Backups are still an important tool in our constant need for data security and protection against loss. If you have a well-rehearsed Disaster Recovery Plan, then you're almost certainly taking backups and moving these off-site somewhere. However, how can you ensure that in the worst case scenario — the backup goes missing, either by accident or through a deliberate act — that the data contained within it is secure?

If you have created an encrypted file, as discussed above, then that data will remain encrypted on the save, as Reality does not decrypt the data to save it. The one exception to this is a `T-DUMP` whereby the data is first unencrypted before being dumped to the media.

If you want to create a fully encrypted backup of all the data, including the `T-DUMP` format, then you must first define an encrypted physical or pseudo-tape drive. You can either do this in the database `config` file so it's defined permanently, or you can define it as a one-off using the `T-DEVICE` verb. The device configuration is the same whether it's in the database `config` file or via `T-DEVICE`.

The syntax is:

```
T-DEVICE { unit { device path
{device type}}}
```

If the `device type` is unspecified, but unit and `device path` are both specified, then it defaults to a tape image (pseudo-tape).

To add encryption also to the tape definition, we also need to append `:kencyrptionkey` where the `encryptionkey` is as we defined above. So an example, using our payroll encryption key to a Reality tape image, would be:

```
T-DEVICE 1 c:/temp/mytape.
rti:kpayroll
```

This specifies tape unit 1 using the file `c:/temp/mytape.rti` and the encryption key referenced by `payroll`. Now we can assign ourselves to this tape unit and begin our backup procedure. All the data, regardless of whether or not it's encrypted at rest, will be encrypted on the backup using the supplied encryption key.

Note: If you wish to restore the encrypted data, then you must first restore the encryption key using `LOAD-REK` as described above, if it's not already present for the database that you are loading into.

This is just a brief discussion of Reality's security model and how encryption fits into this. We have discussed how to create an encryption key, how to use this to create an encrypted file, how to encrypt individual fields of data and finally, how to secure your off-site backups by creating a fully encrypted save. Want to know more? Contact your local representative or go to www.northgate-is.com/reality or check out the online manuals and use the index to find more on encryption. IS

ACCOUNTING

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

COMPLIANCE

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

CONSULTING

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

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

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

Modern MultiValue, LLC
www.ModernMultiValue.com | info@ModernMultiValue.com

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

Precision Solutions
www.precisonline.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/

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

You Could be an Ostrich Farmer!

BY CHARLES BAROUCH

For those who don't know, the young lady in the pilot's seat is my wife and co-editor, Syd Barouch. No, she isn't a pilot. She has, however, been a day manager at a mixed martial arts school, worked in childcare, painted hats for a charitable organization, and raised two kids. Why is that the lead-in for a Business Tech article? Keep reading.

I used to work for a man named Eric. He had an excessively well-maintained computer room. When I asked him why he put the time into it, he said, "I used to be rented, with the audio equipment, when people like Rod Stewart played arenas and concert halls. To me, a cable is something that holds a speaker over the heads of thousands of attendees." We bring these experiences — the other jobs and other lives we live — into our work. Sometimes for the better, and sometimes, not.

Fly on the Wall

Not too long ago, I was on a flight during which I sat next to a recent retiree, and we fell into conversation. She'd done software quality assurance during much of her career. The encounter could have stopped at "My daughter mentors women in technology. Now that I'm retired, I'm helping her with that." But under the conversation about tech being a boys' club and the details about her writing for her daughter's business, lurked a long, rich vein of common topics between us.

Knowing the secondary skills and off-topic experience of the people you work with can often make communicating with them a step closer to effortless.

I learned a lot about a company I had briefly done training for — she'd worked with them extensively — and I picked up a new perspective on non-MultiValue software. We may also have set a record for the number of times Rational Clearcase came up as a point of reference outside of a conference room. It was a random conversation on a plane that will improve my job performance.

Attributes on the Rocks

Knowing the secondary skills and off-topic experiences of the people you work with can often make communicating with them easier. For example, I run into a lot of IT people who are also musicians. This is not my skill set. I have G-clef envy. I admit it. I can't even keep a radio in tune.

Talking about my appreciation for the patience, effort, and dedication that their playing requires allows me to pay them respect. Mutual respect and genuine interest in what they know smooths the way for other — on-topic — conversations.

Gaming the System

Like music, board and card games are referred to as 'play' even when you apply rigor to them. I have friends who are brilliant analysts, but have never written or reviewed code. They use their analytics in their hobby. When you hear things like, "a move that doesn't cost you any resources still costs you a move, and this game only allows x number of moves," you know that you are in the presence of people who don't play around when it comes to games.

This is where I use my on-topic-for-work skills in a social setting. What I do for a living allows me to participate effectively in conversations about game mechanics. Net result: I view hobbies differently when hiring people.

Aboard the Train

When I train people to work in MultiValue, I find that students who posess either an accounting background or a Cobol background tend to make excellent MultiValue programmers. Bonus points for those who have both.

These professional experience sets lend themselves to developing disciplined habits. Building on that, when teaching the – ahem – frequently less disciplined world of MultiValue, helps them connect to our sort of work in a way that leads them toward a more successful path. Linking our other lives to our current work allows us to be better at our jobs.

One of my favorite training stories is one where I was the student, not the teacher. I was brushing up on a language I had picked up on my own. In service of that goal, I took my first formal class even though I was already working professionally in that field.

The teacher was (A) very good at teaching, and (B) not very good at the subject. He didn't have a grasp of several core commands. He didn't understand efficiency. Here's what I learned: He used very few commands over-all, his code was very easy to follow as a result, and it worked. In a classroom setting, it was a brilliant approach. I changed my teaching methods to emphasize smaller code vocabularies and only introduce the better (i.e. faster, more efficient) methods after showing the simpler ones.

This is the corollary to my previous point. He didn't do this because it was a better teaching approach. He simply brought his lack of skill as a programmer into the classroom and made it work. I've followed his postings, so I know this wasn't a clever ploy. The old saying, "those who can't do, teach" is not intended to be a compliment, but this was the unusual case of bad-at-the-job but good-at-teaching.

You Could be an Ostrich Farmer

The title comes from a conversation that occurred at a (non-Spectrum) conference. The lunch area was set up to encourage people to sit with strangers in order to maximize the social net-working aspects of the event. Three of us ended up at a table, set for ten, together. One gentleman declined to talk at all; he was content just to listen. The other man and I engaged in a robust exchange of opinions about the merit, or lack thereof, of various technologies. Finally, I said, “Enough business talk. What do you do when you’re not doing this?”

“I’m an ostrich farmer,” he said.

I looked over at the silent partner in our discussion and decided to have a little fun. Implying I knew the newly-revealed ostrich farmer better than I did, I turned to our third wheel and quipped, “And he lives in a tiny apartment.”

That’s a Wrap

For those of you who were hoping this article would teach you how to become an ostrich farmer, I suggest you strike up some random conversations with strangers and co-workers. There might be someone near you right now with a deep background in the ostrich or emu industry. You’ll never know until you ask.

Send your best ostrich farming stories to editor@intl-spectrum.com

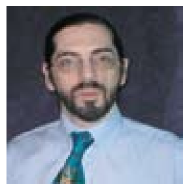

CHARLES BAROUCH is the CTO of HDWP, Inc. and the Publisher at HDWPbooks. You can read his writing in International Spectrum, Theme-Thology, Novo Pulp, PerehelionSF, and the Interrogative series, which begins with Tiago and the Masterless.