PLUS!
Preparing
for Business
Intelligence

The
Art
and
Science
of
Training

INTERNATIONAL Spectrum®

THE MULTIVALUE TECHNOLOGY MAGAZINE

DEPARTMENTS

FEATURES | SEPTEMBER/OCTOBER 2008

From the Inside

BY NATHAN RECTOR

Probably eight out of ten of you aren't using your e-mail in your business correctly.

I know that is very harsh thing to say, but it's true. Because we are IT people and not salespeople, we tend to forget a section of our business and specifically how e-mail is used by that section of business.

The business has focused so much on how to send e-mails from their MultiValue systems, they have forgotten about handling incoming e-mails from their customers and vendors, as well as the outgoing e-mails from their salespeople, accounting departments, and HR departments.

Have I dropped enough hints yet?

Well, look at your e-mail as a database. It is a database of correspondence between two or more people. So what do your salespeople do with your company's application (CRM or MRP system), when they get an e-mail from a customer? If you are lucky, they will add a note to the system regarding what the e-mail is about.

If you aren't that lucky, they will do nothing and you just lost audit and tracking information that may be key to your business. To make matters worse, many company's supply their e-mails to their employees using POP3, which means the e-mail can become distributed and deleted outside the corporate controls.

Have you caught on to the hole in your business systems yet?

In many companies, e-mail is treated just like a phone call. When it comes in, someone reads it, responds to it, and deletes it. And then, if they think it's important, might make a short note in the company's CRM or ERP system of what's going on.

Well, there is no need for this. Your corporate e-mail systems should be attached to your business systems so that when an e-mail is sent or received, there is a note added to the CRM or ERP that the e-mail was received, opened, responded to, and/or sent. If you want to get more advanced, you can even add the text of the e-mail into your systems, all without the salesperson, production manager, or HR person having to do a thing.

Every e-mail has an address. Our databases have e-mail addresses attached to customer and vendor accounts. There is no technical reason why we can't capture this information on incoming and outgoing e-mail, and process the information as needed — even if someone is using Outlook or Thunderbird as their client.

I would encourage you to take a long, close look at your e-mail systems and start including them into your business systems. You'll find a whole new rich set of data that you never knew you had.

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

BY CHARLES BAROUCH

The Art and Science of Training
Part 1

I t's training. It's boring. The whole class is snoring.

So, there I am, sitting with my partner Bob Coltun, listening to this brilliant instructor — who's name I will withhold — and we are looking at the exit every six seconds. Please understand that we both believe strongly in training. We both have reason to be very interested in the material being presented. However, the instructor doesn't possess the teaching gene. He just can't communicate his knowledge.

So, what makes a good trainer? Let's start with what this unnamed instructor did right. He had something of value to communicate. His topic was of interest and he knew a lot about it. Additionally, as we could tell by the way he kept looking down at the cards in front of him, he was prepared.

Where did it go wrong? He spoke in a monotone. He focused his eyes, and therefore his face and his voice, at the note cards; so we had trouble hearing him. His entire approach effectively communicated his fear of public speaking, but never communicated his passion, which was so great that it drove him to present the topic despite his fear. How do we do better when we train others?

Topic-al Depression

When I speak at Spectrum and other conferences, I have the luxury of suggesting topics I want to cover. In most training situations you have to teach what you've been asked to teach. Since some topics are inherently more exciting than others, a good trainer looks for the hook. For example, when talking about business workflow, you could do a session called "Working Smarter, Not Harder - Efficiencies in Business Workflow Seen Through Five Case Studies" or you could do a session called "Five Ways to Leave at Five." The title underlines the first rule: *You can't train them if they won't show up.*

Remember that a title is a promise. If I do attend a session with the first title, I have to expect that specific examples — five case studies — will be covered. I have to expect that trainer will tie these examples together into a coherent argument in favor of exacting specific changes. If I attend a session with the second title, I expect a vastly less formal presentation, with an emphasis on practical advice. In either case, the trainer has to

understand our second rule: *Keep your promises if you want them to keep coming back.*

Content Under Pressure

When my friend, Jeff Downing, adds a class about roses to the schedule at The New York Botanical Garden's Continuing Education offerings, he thinks about the audience. He can add a class for the occasional gardener, for the commercial landscaper, for the botanist, and so on. Obviously one class will not serve all of these people equally. He needs to make that decision and communicate it to the students, the instructor, and to his staff.

With IT training and business training, we can safely assume certain things about our audience. Take it as a given that they are under pressure. If I train someone in web integration, the likelihood that they have a web integration project planned is very high. If I train someone on database administration, they probably already have a job doing that and are under pressure to learn and apply the training. Knowing our audience helps us to select a title and content that will connect the person to the information. Rule 3: *Know your audience.*

The Hook

When I did Train-the-Trainer for IBM, I started my presentation like this; "Hi, I'm Chuck Barouch. I don't have any slides, so you'll have to take your own notes. This should be over pretty quick." As you might have guessed, this is an awful opening. It contains several of the sins associated with bad training.

This was done to prove a point. Picking apart that bad opening with the attendees gave me an excellent way to introduce the topic of how to train.

> "
> *If you can't find anything in common with your audience, expect the training to go badly. A common frame of reference is critical.*
> "

Until you open your mouth, you have some earned credits with the audience. They assume you belong up at the front of the room. Any opening that doesn't blow that advantage is a plus. Any opening which builds on that advantage will make the rest of the session easier.

So, what are the key sins in my intentionally bad opening? I started by telling them my name but not my credentials. They deserve to know if I deserve their attention. Of course, following up with an apology for being unprepared doesn't inspire a great deal of confidence. Worst of all, I gave the audience no clue as to what my topic or tome would be. People want to know what to expect.

My hook usually involves establishing empathy. People like to know that you have had common experiences. It makes you one of their tribe, a part of the group. Comedians know this. Think about the number of comics who say things like "have you ever noticed" or "we've all had that moment when" as a way of including the audience. Please be aware that faking the hook is a great way to offend the room. If I stand in front of a bunch of Doctors, I can start with a hook about medical testing because I've done oversight in that area, but I won't pretend I perform surgery just to make an inclusive observation.

If you can't find anything in common with your audience, expect the training to go badly. A common frame of reference is critical. One thing which may help is to spend a day in their shoes. Obviously, I can't become a surgeon for a day, but I can sit at a customer service post and take calls before I train them in the new software. I can walk the floor with a shipping document or ride along with a distribution delivery. This all goes back to Rule 3: *Know your audience.*

The Line

Some people think that humor is essential to all presentations. This is not true. Here are some guidelines for deciding if humor is appropriate:

The Five Rules of Training

1. You Can't Train Them if They Won't Show Up.

2. Keep Your Promises if You Want Them To Keep Coming Back.

3. Know Your Audience

4. If Laughter is the Best Medicine, Then Use It in Controlled Doses Only.

5. Have Beginning, a Middle, and an End.

BY SHANNON STOLTZ

The Top 5

Customer Service Mistakes IT Professionals Make

We are all human; we all make mistakes. One of the reasons I write and teach on customer service is because I've learned some lessons the hard way; and I've watched many others make the same mistakes over and over again. Normally, I prefer to put a positive spin on life and mistakes, but sometimes it's necessary to look at a list of don'ts.

So today, let's take a look at the top five customer service mistakes IT professionals make and see if any resonate with you. I know I'm guilty of making these mistakes, even today as I teach on the topic. We're human.

The point is not to say "I'm already there and I'm great!" or "I'm never going to be there so I shouldn't even try." Instead, the point is to identify an area to grow in and improve your skills. And as you do that, you will see easier relationships with your customers and more potential for job growth or new opportunities.

Remember, a customer is not just those revenue-generating customers, but also anyone who uses your services, including co-workers and managers. So having said that, let's look at the top five mistakes IT professionals make in the area of customer service.

It's Not My Job

Mistake number one: Refusing to help someone who asks you a question that doesn't pertain to your area of expertise.

Let's face it, our customers don't usually understand how their computer works or that there are differences between applications and databases and operating systems, or even hardware and networks. Our customers usually don't know how the IT department is structured and who to go to ask what. They just know they have a problem or something that needs to be done and you know something about computers. When they come ask you a question, they just want help or to be connected to someone who can help. They are looking to you to get them on the right path.

So, if someone comes and asks you a question that isn't exactly your responsibility, remember that clause at the bottom of your job description that says, "And all other duties, as assigned." No, you are not expected to solve their problem. But you are expected to physically connect them to the person who can help them.

Have you ever called a customer service line or technical support and told your story to one person, got put on hold and moved to another person, had to re-tell your story, worked through some issues, got passed to yet another person, and had to start all over with your story and everything that's been done up to that point? If you have, you know this is a frustrating, time consuming experience that tests the patience of the most patient person.

So, the next lesson here is, when you do point someone in the right direction, help them out and provide a real introduction and explanation to the person you are handing them off to. And then, circle around later and just make sure that the person got their answer solved and isn't just left hanging — which brings us to mistake number two.

Leaving Customers Hanging

Mistake number two: Not closing the loop with the customer, making sure they have the information or solution they need.

When I do customer satisfaction and employee satisfaction surveys for my clients, this item comes up frequently in the comments. And really it is multi-fold. The first is the scenario I just mentioned. When you hook someone up with someone else for help, just do a touch point later on, "were you able to get that problem resolved?" or "was Fred able to help you with that report request?" This is just a way to let the customer know that you respect that they had an issue and want to make sure they found a solution. It is also a way to build a relationship with that customer, which then makes everyone's life easier later on.

The second part of this is when you say you are doing to do something, to do it. Or at least communicate the status when the deadline is rapidly approaching. This is about communication, whether it be via e-mail, phone, or in person passing through a hallway. Follow up and communicate with the customer.

Recently I was supposed to mock up a custom report for a programmer by a certain day. That day came and there was no way I was going to get it done. I could have just let it be and waited until he called frustrated asking "weren't you supposed to send me this?" Instead, he appreciated that I e-mailed and said, "Listen I know I am supposed to get this to you, but I'm not able to do it today. Is tomorrow okay?"

Reset your expectations. Don't leave them hanging and don't assume that you can just then drop it and somebody else will pick it up. Do what you say you're going to do and if you can't, re-set expectations. But don't just leave them hanging.

Now, I know that for some this is a troublesome area, because some may say this can be akin to admitting a mistake or not being perfect or not being able to do it all. But I guarantee, your customer would rather know that a delay is coming and than being sitting there on the deadline day wondering where it is.

The point is, keep the chains of communication open. It makes a difference. So let's look at mistake number three.

> "
> *But despite the fact that we can program computers and make them do what we want, we are not gods, and we certainly can't know everything there is to know.*
> "

Guessing

Mistake number three: Guessing when you really don't know the answer.

When someone asks a question or mentions a topic we really don't know a lot about, it is sometimes tempting to "fake it" or to guess, instead of finding the correct answer or hooking up the customer with someone who knows. For some it can be difficult for the ego – they want people to think they know everything. But despite the fact that we can program computers and make them do what we want, we are not gods, and we certainly can't know everything there is to know.

It is frustrating for a customer to be working with someone and then later find out that if they had gone to someone else they could have the problem solved much quicker. The reality of our work is that performance and time matter. If we can't provide a real answer or real solution, then we need to say "I don't know, but let me find out for you" or "I'm not sure, but let's go talk to so-and-so". (Again, hook that customer up with the expert.) It is through finding answers for customers that many gurus become gurus.

Now, don't get me wrong, I'm all for experience-based, intelligent guessing in a trouble shooting environment or for establishing a starting place to start looking into something. But if you are guessing, couch it with a "Listen I'm guessing here but let's try this..." Then, if it works, you are a genius; and if it doesn't, the customer knows it was guess, a possibility, and not fact. No harm done. But stating a guess as a fact can backfire. It's just better to keep those expectations as honest and up-front as you can.

This mistake of guessing when you really don't know the answer is often related to our next mistake.

Too Much or Too Little Explanation

Mistake number four: Providing too much information or too little information in your explanations.

Ever asked someone, "how are you" and got a long detailed answer about everything they are going through right now? It's a little frightening, isn't it? Here you are being friendly and expecting the usually "good" or "fine" answer, and they spend five minutes answering the question. Well the same thing happens in the IT world.

Most people don't need a long detailed answer on a technical issue. They need the *Reader's Digest* version, and sometimes just something very simple. Now I admit this is a tough one; it requires you to know your customer a bit and to learn to read when you are giving too much information. If you pay attention, you'll see the visual clues – glazed eyes, slowly backing away, shifting from side to side, the annoyed or "deer in the headlights" look that passes over their face.

Now, the tough part of this, is the opposite. Some people do want or need more of an explanation. By not giving enough information, you can be hampering that person's ability to do their job. When I'm doing business analysis

UniData Basic

via CODE PROFILING

PART 3

BY KEVIN KING

Throughout this series we've been looking at how easy it is to leverage UniData's built-in profiling features to improve the performance of our pro gramming. In short, compile with –G, run with –G, and UniData rewards you with a bounty of useful optimization information.

We've seen how UniData generates two different profiles, one based on elapsed CPU time (profile.*pid*) and one based on elapsed clock time (profile.elapse.*pid*). So far we've been focusing on the elapsed clock time profile (fig. 2) to help us understand the performance of our sample code. But does that imply that the CPU statistics are any less valuable?

Here's the issue: One processor can serve many different masters simultaneously. From a user perspective, let's say Harry and Sally are the only users on the machine. Furthermore, for demonstration purposes, let's say (technically incorrectly) that there are only two programs running on that box, one for Harry and one for Sally. Both programs were started at basically the same time, and let's say — again, building this hyper-unrealistic example — both finished on the same second, exactly two minutes later.

Looking at the elapsed clock time profile for these programs, we'd see that both of them took the same two minute window, so does that mean the two programs are equally efficient? Not necessarily. While two minutes may have clicked off on the clock, the single CPU on that machine bounced thousands of times between running Harry's program and running Sally's program. The efficiency of each program might then be measured by comparing the elapsed CPU time for each session.

But even this approach has its flaws. If Harry's program has more I/O than Sally's, his program might run in the same two minute window and yet reflect a more efficient CPU utilization overall. In other words, just because a program is running slow doesn't necessarily mean that the CPU itself is being bogged down. With that in mind, how exactly are we expected to make sense of these CPU statistics?

Figure 1 shows a small program created on UniData 7.1/Windows that repeats an otherwise meaningless calculation a half a million times. Figures 2 and 3 show the elapsed and CPU profiles for a run of this program. If we compare the cumulative seconds between these two profiles we could say that for each of the seconds that this program ran, the CPU load per second could be calculated as 0.009375, or .3 total CPU seconds divided by 32 elapsed seconds.

Figure 4 shows a variation on this program; here's a slightly less involved calculation, again run a half million times. Figures 5 and 6 show the elapsed and

```
FOR TIGHT.LOOP = 1 TO 500000
    ANS = ICONV(OCONV(RND(100)*71/3,'MR2'),'MR2')
    PRINT ANS
NEXT TIGHT.LOOP
```

Fig. 1 Sample UniData 7.1/Windows program

```
%time   cumsecs         seconds         calls  name

100.0   32.00           32.00           1    TRAINPROGS\_HOT.CPU


                                        called/total     parents
index  %time   self            descendents  called+self      name      index
                                        called/total     children

                                                         <spontaneous>
[1]    100.0  32.00            0.00              1        TRAINPROGS\_HOT.CPU [1]

---------------------------------------------
```

Fig. 2 Elapsed profile of sample program

```
%time cumsecs seconds    calls name

100.0    0.03    0.03            1 TRAINPROGS\_HOT.CPU


                              called/total        parents
index %time   self   descendents called+self      name        index
                              called/total        children

                                                  <spontaneous>
[1]    100.0  0.03            0.00    1            TRAINPROGS\_HOT.CPU [1]

---------------------------------------------
```

Fig. 3 CPU profile of sample program

```
FOR TIGHT.LOOP = 1 TO 500000
    ANS = RND(100) * 71 / 3
    PRINT ANS
NEXT TIGHT.LOOP
```

Fig. 4 Variation of sample UniData 7.1/Windows program

```
%time cumsecs seconds    calls name

100.0   34.00   34.00        1 TRAINPROGS\_WARM.CPU


                              called/total        parents
index %time   self   descendents called+self   name   index
                              called/total        children

                                                  <spontaneous>
[1]    100.0  34.00            0.00        1        TRAINPROGS\_WARM.CPU [1]

---------------------------------------------
```

Fig. 5 Elapsed profile of sample program in figure 4

```
%time cumsecs seconds    calls name

100.0    0.03    0.03        1 TRAINPROGS\_WARM.CPU


                              called/total        parents
index %time   self   descendents called+self   name   index
                              called/total        children

                                                  <spontaneous>
[1]    100.0    0.03    0.00            1        TRAINPROGS\_WARM.CPU [1]

---------------------------------------------
```

Fig. 6 CPU profile of sample program in figure 4

Fig. 7 Windows Task Manager

OPTIMIZING UNIDATA BASIC VIA CODE PROFILING, PART 3

Continued from page 10

CPU profiles for a run of this program. Notice how the results are basically the same. (The second program ran a little longer than the first, but I believe that may be due to the fact that the numbers being output were slightly longer.)

What you may find interesting about this is that while both programs show a miniscule CPU load, according to the Windows Task Manager (fig. 7), each of these programs completely pegged the CPU during the run. So could that mean that the CPU statistics are meaningless on Windows? Possibly. But before we go making those claims let's take a look at the same two programs running under AIX.

Figures 8 and 9 show the profiles for our first program on a heavily loaded AIX box. Not only does it take a significantly longer time to run in that environment, dividing the clock time by the elapsed time we see a more representative 0.0203 load. Furthermore the AIX "topas" reported only a .3% CPU load
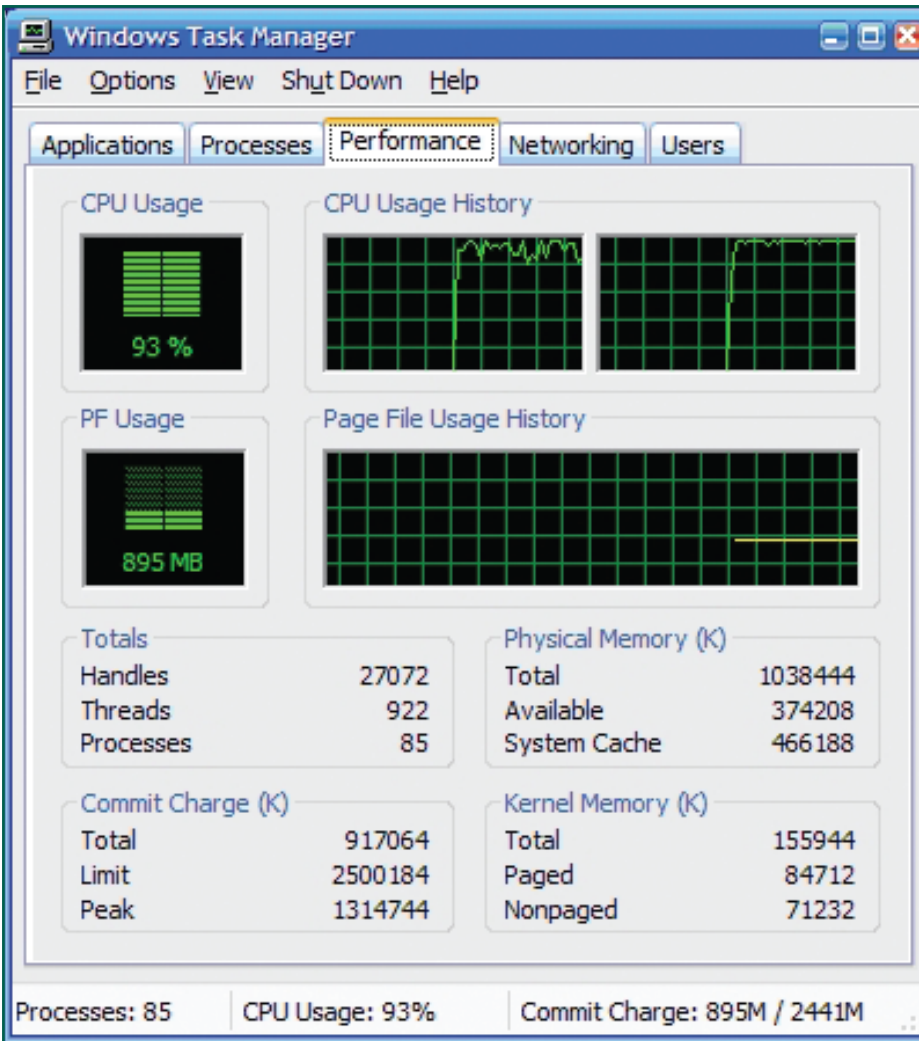
```
%time cumsecs seconds    calls name

100.0  389.00  389.00        1 EXEC/DEVPROGS/_HOT.CPU



                          called/total       parents
index %time   self   descendents called+self    name    index
                          called/total       children


                                          <spontaneous>
[1]     100.0 389.00       0.00        1      EXEC/DEVPROGS/_HOT.CPU [1]

-----------------------------------------------
```

Fig. 8 Elapsed profile of figure 1 program on a loaded AIX box

```
%time cumsecs seconds    calls name

100.0    7.90    7.90        1 EXEC/DEVPROGS/_HOT.CPU



                          called/total       parents
index %time   self   descendents called+self    name     index
                          called/total       children


                                          <spontaneous>
[1]    100.0  7.90         0.00        1      EXEC/DEVPROGS/_HOT.CPU [1]

-----------------------------------------------
```

Fig. 9 CPU profile of figure 1 program on a loaded AIX box

Phantoms

Do More Than

Go Bump In The Night

BY CANDI HART

Today's reality consists of the sharing of information existing in worlds other than only MultiValue. Our systems need be able to determine when another system has data available for us without relying on some vague procedure or sporadically executed process. Phantoms or background processes, though not seen, *are* reliable and don't require a physical presence to check for new data. Even though in my last article on KPI's I replaced a phantom process with a trigger, I have found phantom processes to be a valuable tool, especially when interfacing with other systems.

The Phantom When Restricted To Its Own Environment

Phantoms or background processes have been available on MultiValue operating systems since they were proprietary operating systems and everyone used terminals.

They were used most often to set off a report that took a long time to run and then release the terminal so the user could continue to work. This, you could say, was the recognition in the MultiValue community of the need to multitask long before Windows and PCs. Over the years, as PCs began to replace terminals and users could have several windows open, I have still found many applications that work best in the background without physical intervention.

The initial use of phantoms was to kick off some process from TCL and the phantom would log off when the process completed. By creating a program that contains a loop and a sleep command, then running that program as a phantom, the phantom becomes more than just one procedure running without a physical terminal, but a continuously running process that can automatically do regularly scheduled functions (fig. 1).

Some example application modules where a continuously running phantom is useful:

```
001 * FILE.MONITOR
002 * This program  should be started as a phantom.
003 * It will run continuously until the control item is set to 'stop'
004 *
005 OPEN 'CONTROL' TO CONTROL.F ELSE STOP
006 OPEN 'IMPORT.DIRECTORY' TO IMPORT.F ELSE STOP
007 *
008 LOOP
009     READ CONTROL.REC FROM CONTROL.F,"%CONTROL%" ELSE CONTROL.REC="STOP"
010     STATUS = CONTROL.REC<1>
011 UNTIL STATUS = "STOP" DO
012     SELECT IMPORT.F
013     MORE = 1
014     LOOP
015         READNEXT TXTREC.ID ELSE MORE = 0
016     WHILE MORE DO
017         CALL FILE.PROCESS.SUBROUTINE(TEXTREC.ID)
018     REPEAT
019     SLEEP 300
020 REPEAT
021 *
022 END
```

Fig. 1 Basic program to create a continuously running polling program

- As a night process to run regular reports. (Microsoft has the Windows scheduler for Windows tasks that need to run regularly without operator intervention.) Within a MultiValue application, a simple scheduled reports process might be used to create a list of reports that need to be run each night. They can be cataloged programs or procs that can run from TCL. Then a Basic program is written that wakes up each night at a scheduled time — say after the file save has run — that reads the list and does an execute of each report. When it is complete the program would sleep until the next night.

- The nighttime can be a great time to do purges of files that do not contain critical information but seem to build up — error logs for instance, or hold files that are more than 30 days old.

- Printing bar code labels in the warehouse when purchase orders are received can take a long time due to the speed of the label printer. Rather than tie up the receiving process, the part numbers received, quantity, and PO number can be written to a label file, and a phantom can start the labels printing whenever any items showed up in that file.

- In another instance, the night crew was using bar code scanners to stage inventory for the morning shipments. Several files needed to be updated when the product was pulled, and some of those files might be locked (due to the filesave or other users accessing the inventory files). This left the forklift operator waiting for the update to release. All validations of the data was done real time, but it made sense to write the transaction to a transaction file and do the file updates later by a background process. That phantom was waking up every five minutes to process transactions throughout the day, but during the night it could skip the hours during the filesave.

Though this may seem like going backwards to batch processing instead of real time, the sleep time can be minimal and this is more like the modular design concept that is needed for the new world. For example, web design where we need to see the capturing of data as separate from the processing of that data. In the last example above, it is possible for the company to upgrade to any of the new graphical scanners and data collection devices as the captured data is simply placed in a file for the MultiValue application to process. This brings us to turning phantoms loose in the world at large.

Phantoms At Work In The Larger Workplace

By creating a file pointer to an external file/directory on the system, the phantom loop can be used to watch for data from other sources, making the process extremely powerful. Following are some instances where I have used this design to interface with other systems:

- Receiving EDI orders. The EDI software ran on a PC in the office and when the operator did a Get and Export the new PO was written to a Windows directory

OPTIMIZING UNIDATA BASIC VIA CODE PROFILING, PART 3

Continued from page 12

for this program, hardly pegging the box like it did on Windows.

Figures 10 and 11 show the same results for our second program on the AIX box. Again, "topas" reported a mere .2-.3% CPU load. Yet, applying the same calculation of CPU time divided by clock time, we can see that this program with the simplified calculation had only a per-second CPU load of 0.0076 – a small but noticeable improvement.

For one last test, I took these same programs and commented out the PRINT statement. On Windows, this reduced both clock and CPU times (and hence overall indicated load) to next to nothing. On AIX, the elapsed time was reduced to next to nothing yet the CPU load per second indicated basically the same as the previous test. This confirmed — at least to me — that the methodology that Windows uses for scheduling the CPU is incompatible with the CPU profiling statistics generated by UniData. On AIX, however, where the CPU appears to be scheduled differently, watching the CPU statistics can give us some good measures of improvement (or not!) as we work through our optimization strategy.

So what exactly increases CPU time? Oddly enough, through these tests it would appear that PRINT really bogs the CPU. The more printing, the heavier the CPU load. This could explain why running reports in the middle of the day tends to cause the machine to just feel slower and more bogged down overall. And all this time we thought it was SELECTs? (Yes, it is likely that SELECTs are still a culprit, but I found it surprising that PRINTs had such an impact!)

Another thing I've noticed — reflected briefly in a comparison between figures 9 and 11 — is that string manipulation (especially substrings and masking) and ICONV/OCONV can have a noticeable impact on CPU load. Not entirely certain why this is the case, but considering that with UniData there's usually a lot of calculations in play to support all this easy string manipulation, it's not completely a surprise.

So now that we have some ideas of what's bogging the CPU, what can we do about it? For starters, maybe we could minimize redundant and otherwise unnecessarily redundant ICONV and OCONV usage. I've seen countless programs that embed one inside of the other, when in reality the calculation might be manageable with a single code, as in figure 12. Or, "take the value, shift the decimal point to the right two places, then shift the decimal

```
%time   cumsecs        seconds        calls  name

100.0   864.00         864.00             1  EXEC/DEVPROGS/_WARM.CPU


                                     called/total          parents
index   %time          self     descendents  called+self    name          index
                                     called/total          children

                                                       <spontaneous>
[1]     100.0          864.00              0.00          1   EXEC/DEVPROGS/_WARM.CPU [1]

---------------------------------------------
```

Fig. 10 Elapsed profile for figure 4 program on a loaded AIX box

```
%time cumsecs seconds    calls name

100.0   6.64    6.64          1 EXEC/DEVPROGS/_WARM.CPU


                               called/total          parents
index %time    self    descendents called+self    name     index
                               called/total          children

                                               <spontaneous>
[1]     100.0   6.64       0.00       1        EXEC/DEVPROGS/_WARM.CPU [1]

---------------------------------------------
```

Fig. 11 CPU profile of figure 4 program on a loaded AIX box

```
 ANS = OCONV(ICONV(VALUE,'MR2'),'MR2')
```

Fig. 12 Decimal movement sample code

Hostile environment?

...not if you
swap data using Reality

Reality is the supremely interoperable database. It crosses environments with ease, importing and exporting securely across any network or over the web using Web Services. With Reality, virtually all data becomes transparent: Relational or MultiValue data, XML files, CSV files, platform host files, or any ODBC-compliant data.

↑northgate
REALITY

www.northgate-reality.com   reality@northgate-is.com

BY ROSS MORRISSEY

Preparing for
Business Intelligence

Sooner or later, every MultiValue shop will need to rework their data for some analytical reporting project. From the full-blown data warehouse to the humble pivot table, successful analytical reporting projects rely on dimensional modeling. This article will introduce dimensional modeling and describe some of the decision making that go into translating MultiValue transactional data into a basic dimensional model. I've refined this design process with IT and business staff from over one hundred MultiValue shops over the past few years; it's straight forward, reduces risk, and can guide both business and IT users to powerful solutions.

Dimensional Modeling

Dimensional modeling originated with the design of large databases that aggregated grocery store sales for General Mills in the 1960s. It was popularized by Ralph Kimball in his 1996 classic "The Data Warehouse Toolkit" (now in an excellent second edition). Dimensional modeling revolves around facts or measurements and dimensions that identify facts. A typical fact is a count or amount — a quantity or dollar figure. On its own, a number doesn't tell the whole story, e.g. 42. Dimensions — like customer, store, product, and time — provide the context for analysis. This fact-dimension relationship can be captured in a classic data warehousing star schema or a flat file, but the design decisions are the same. A quick example will show the difference between facts and dimensions.

All the examples in this article use the HS.SALES database that ships with UniVerse. Figure 1 shows the first two customer records. In this case a fact record might include the price. Obvious dimensions would be customer, product, and time. A simple RetrieVe statement pulls this information out (fig. 2).

Notice that some customer names are repeated when there are customers with multiple transactions. This (de)normalization or flattening of the data is tightly bound to the dimensional model and will be covered shortly. This flattened data lends itself to analysis with pivot tables or for inclusion in many business intelligence solutions. Some of these solutions will handle the extraction or flattening, but not decisions about which fields to use as facts and dimensions. The remainder of this article will focus on these decisions.

The Dimensional Modeling Process

There are four steps to this process:

1. Formulate a vision statement,

2. Identify the transaction granularity,

3. Select a primary source file, and

4. Choose the dimensions and facts.

```
Customer ID....           1
Contact Name... Mr. Samuel Smith
Company Name... Better Beer, Inc.
Street Address.10 Commercial St.
City.......... Concord
State......... NH
State name..... New Hampshire
Zip........... 02131
Telephone...... (603)555-3212

Product Product Description. Serial# Date Purchased Date paid List price Price.. Discount
M2000   Moderate duty,      501278         01/07/91  01/28/91     $4,490 $4,200      6.5
        monochrome copier

Customer ID....          10
Contact Name... Dr. Andrew McCaig
Company Name... HGT Dental Center
Street Address. 999 Hill Road
City.......... Brattleboro
State......... VT
State name..... Vermont
Zip........... 03356
Telephone...... (802)555-6534

Product Product Description. Serial# Date Purchased Date paid List price Price.. Discount
M1000   Low cost, entry     203510         01/28/91  02/14/91     $1,990 $1,990      0.0
        level, light duty,
        monochrome copier
M1000   Low cost, entry     203600         01/29/91  02/28/91     $1,990 $1,900      4.5
        level, light duty,
        monochrome copier
C2000   Moderate duty, entry 600791        01/30/91               $6,890 $6,500      5.7
        level, color copier
```

Fig. 1 Example: first two customer records

```
SORT CUSTOMER WITH PRODID NE "" CUSTID COMPANY STATE BY-EXP PRODID NE ""
    →      PRODID BUY_DATE PRICE ID-SUPP
Customer ID     Company Name........    State     Product     Date Purchased      Price..

        10      HGT Dental Center       VT        C2000          01/30/91         $6,500
         2      Fast Copy Center        MA        C2000          01/08/91         $6,600
         4      Fast Copy Center        MA        C3000          01/09/91        $16,500
        10      HGT Dental Center       VT        M1000          01/28/91         $1,990
        10      HGT Dental Center       VT        M1000          01/29/91         $1,900
         5      Ocean State Fish        RI        M1000          01/14/91         $1,900
                Company
         5      Ocean State Fish        RI        M1000          02/15/91         $1,900
                Company
         1      Better Beer, Inc.       NH        M2000          01/07/91         $4,200
         3      Fast Copy Center        MA        M2000          01/08/91         $4,250
         7      Central Hospital        NH        M2000          01/20/91         $4,490
         2      Fast Copy Center        MA        M3000          01/08/91        $12,000
         8      Copies, Inc.            MA        M3000          01/21/91        $12,000
         7      Central Hospital        NH        S2000          01/20/91           $990
         2      Fast Copy Center        MA        S3000          01/22/91           $900
         8      Copies, Inc.            MA        S3000          01/21/91           $900

15 records listed.
```

Fig. 2 Retrieve statement

NEWSMAKERS

Lyppard Australia Finds Efficient, Cost-Effective Solution with Uniware

Lyppard Australia Limited is a leading national veterinary supplies wholesaler to the Australian and export market. With branches in Melbourne, Sydney, Brisbane, Townsville, Perth, Adelaide, and a dedicated export department, they satisfy requests for product to all corners of the globe.

 "We have a fairly dynamic operation, with over 23,000 product lines and in excess of 500 orders a day and until recently, a headache of a whole lot of paperwork," said Brian Oakley, managing director.

All Lyppard branches warehouse products in accordance with the manufacturer's storage specifications, with optimum stock holdings in all branches to guarantee supply. Strict stock control is maintained to ensure fast and accurate order fulfillment.

The Challenges For Lyppard Australia

In late 2005, Provet Supplies (Vic) Ltd, keen to expand their Victorian and South Australian distribution activity to the national stage, developed the opportunity to acquire the national business of a prominent competitor Lyppard Holdings Ltd. The new merged business, renamed as Lyppard Australia Limited, commenced trading in October 2005.

Managing Director Brian Oakley, while aware of the many challenges to be overcome, remained focused on the

enormous opportunities this presented to the business, the shareholders, and the staff. Intuitively, he recognized that common business process and data in all locations would enable the speediest introduction of efficiencies in all sales and warehouse locations while delivering economies of scale to the national administration enabling the accurate measurement of performance at all levels.

Not wanting to inherit the large IT overheads of the acquired business, Brian made the decision to outsource his entire IT requirement to a trusted single source provider to manage the integration of the business onto a common, efficient, and stable platform.

Uniware Pty Ltd was selected as this partner. A project team comprised of key stakeholders at Lyppard and Uniware was formed to plan and execute the strategy.

Background To IT Systems

Prior to the acquisition, Provet Supplies (Vic) Ltd was running the business on the Distrib/Pick ERP(Enterprise Resource Planning) system. This system, while deemed to be highly efficient and easy to use by both management and staff, was perceived to have some issues for the moving forward strategy.

Lyppard Holdings Ltd at the time of acquisition was using the SAP system. This system,

with a high cost of ownership, inefficient business processes, and associated hefty telecommunication costs for the national private network, had not enhanced the business. Thus, Brian Oakley dismissed it as a suitable platform to move forward with.

Realizing The Opportunities

Brian Oakley and the project team were committed to providing the shareholders with the quickest possible return on investment and implementing efficient, measurable, and replicable processes in the most cost effective and timely manner.

To meet these commitments it was decided to move all operational and administrative process to the AusVantage ERP suite from Uniware operating on IBM's Unidata Database Management System with key target dates set within the implementation plan. These targets comprised logistical, IT infrastructure, ERP system implementation, and e-commerce enablement.

The most significant step was to have all locations processing on the new Ausvantage ERP suite within three months of merging the businesses. This was achieved with the last branch coming on line in January 2006.

Uniware's IT services organization managed the change from an expensive private network to

a stable and efficient Virtual Private Network reducing overall voice and data costs by 50%.

The cultural challenges and fears of implementing a new ERP system were quickly overcome when, during training sessions, the system users identified the significantly reduced number of steps to manage customer orders to fulfillment making it easier for them to meet their KPIs.

Lyppard Today And Tomorrow

Today, Lyppard operates a highly efficient business, processing over 200,000 order lines each month with a turnover of approximately $150,000,000 per annum and continuing to capture market share.

Approximately 25% of orders are received electronically from the web ordering system, integrated Vet Practice systems, and PDA devices. Online ordering continues to grow steadily. Orders are processed in all locations on the wireless and paperless Ausvantage Warehouse module.

Brian recently commented to his board "we have now squeezed all possible efficiencies within the business and it is now time to take the next step." Lyppard is now seeking to replicate this success with further acquisitions and the systems are all in place to facilitate this corporate agility.

About Uniware

Uniware provides SME's in the Supply chain, Manufacturing, and Services sectors with ERP solutions to manage business processes. With added value services such as IT infrastructure, Web Services, and E-commerce, Uniware is an end to end IT source satisfying our clients' needs to enhance enterprise productivity, mitigate commercial risk, and reduce capital expenditures. For more information about Uniware, contact Uniware Pty Ltd at +61 3 8804 0804 or www.uniware.com.au ■

Ashwood Computer & C-7 Reporting Solutions Join Forces In Strategic Partnership

Ashwood Computer, Inc. and C-7 Reporting Solutions LLC have reached a strategic marketing agreement. The combined sales resources of C-7 and Ashwood will work together to offer their customers a more complete line of products and services.

Rod Owens, president of Ashwood said, "Ashwood's customers will now benefit from C-7 Reporting Solution's expertise with popular business intelligence tools and products including myViewpoint, MITS, and Tantivia Velocity. Additionally, C-7 will now be offering Ashwood's extensive product line including our new Manufacturing Software Applications, our FastBac disk-based backup

toolset, mv.NET, and also, our legacy software conversion services and GUI development expertise. Our partnership and these combined offerings will enable us to provide a much wider range of improved products and services to MultiValue system users nationwide. We really are excited to be joining forces with Larry and C-7 Reporting Solutions."

"We couldn't be happier with this alliance," said Larry Christensen, managing member of C-7 Reporting Solutions. "We have worked with Ashwood for many years and have always been impressed with their top notch products and their highly skilled technical team. They know what it takes to make a user happy. Working together, the combination of products and expertise we bring to the table is going to be very difficult to beat."

About C-7 Reporting Solutions, LLC

C-7 Reporting Solutions, LLC leverages "Best of Breed" solutions and matches their strengths to solve end user and dealer reporting problems. We are a full time reporting solutions company dedicated to our customers reporting requirements. We assist managers wanting to increase sales and profits while reducing costs. Contact C-7 Reporting at 909-349-1736.

About Ashwood Computer, Inc.

Located in Cincinnati, Ohio, Ashwood has been a MultiValue database industry VAR since 1989, providing sales assistance with database servers, Unix, Linux, and Windows operating systems, MV databases, peripheral products, and services and support for their client base nationwide. Ashwood specializes in assisting legacy MultiValue computer systems users with their

conversions and migrations to Open environments, in optimizing systems performance, and by offering a complete end-to-end services solution thereafter. Contact the A-Team today at (513) 563-2800. ■

Mexico-based Ixpan Joins BlueFinity International's Distribution Network

BlueFinity International, a member of the Mpower1 Group of companies, is pleased to add Ixpan, S.A. de C.V. (www.ixpan.com.mx) to its growing network of distributors of the BlueFinity product line. The agreement includes both mv.NET, the world class Microsoft .NET development and deployment product for all the major MultiValue databases, and Reporting Services Data Connector (RSDC) for Microsoft Reporting Services solutions.

mv.NET is a comprehensive solution for accessing MultiValue databases from within Microsoft's .NET environment. It enables software designers to combine the power and flexibility of proven MultiValue technology with the feature rich .NET environment.

RSDC, when used with mv.NET, supports the full feature set of Microsoft's Reporting Services product and allows it to be used against all major MultiValue database platforms with almost no extra work on behalf of the application developer. This allows MultiValue databases to provide sophisticated, real-time report

data using a wide variety of delivery mediums (browser, PDF, rich-client, application-embedded), incorporating all of the reporting features expected by today's demanding report consumer.

Headquartered in Mexico City, IXPAN specializes in systems integration and the development of custom applications in Central and South America with expertise in both Microsoft .NET and MultiValue databases. They work with both end users and VARs to provide custom solutions to their clients.

This agreement is part of an initiative to provide direct sales and support for a growing client base in Central and South America. "BlueFinity's rich product portfolio will allow Ixpan to delve deeper into the largest local and international companies that are already our customers, as well as engage with new prospects," explains José Martín Rodríguez Morales, CEO of Ixpan. "We expect that teaming up our skilled sales channel with BlueFinity's best-in-class products will create a very successful relationship."

This distribution agreement is only the latest step in realizing the market potential of BlueFinity technology in the area. "BlueFinity has been looking to expand it's presence into Central and South America," says Pete Loveless, CEO of BlueFinity International. "We are confident that the combination of BlueFinity's products with Ixpan's local knowledge and comprehensive services will ensure the success of the product range in the region."

For more information about BlueFinity products, visit www.bluefinity.com.

About BlueFinity International

BlueFinity International, a member of the Mpower1 Group of companies (www.mpower1. com) offers its two flagship products — mv.NET and RSDC — to the global MultiValue community. mv.NET allows the full benefits of the .NET service oriented architecture technology to be realized by users of established MultiValue applications; RSDC enables MultiValue developers to utilize the very latest Microsoft Reporting Services technology. For more information, visit www.bluefinity.com. ■

ISCO Industries Keeps Pipeline Full with MITS Discover Business Intelligence

ISCO Industries is one of the largest suppliers of HDPE (high-density polyethylene) pipe and related services in the United States. Founded in 1962 in Louisville, Kentucky, the company distributed a progression of pipe categories over time. In 1975, ISCO began developing its expertise in HDPE pipe, along with the fusion equipment used to install it.

In 2006, ISCO began managing its business with the Activant Prelude enterprise resource planning (ERP) solution. At the same time, the company's geographically widespread sales force consisting of outside reps and a large, centralized order processing unit recognized the need for easy-to-access, up-to-date information on customers and orders.

Data Access Requirements Both Commonplace And Unique

Kim Byram joined ISCO as Director of Information Systems shortly after their decision to purchase Activant Prelude and was charged with overseeing the Prelude implementation. "Our sales force, including about 80 inside and outside reps, needed to be able to get information fast — a customer's order history, current order status, receivables balance, and so on." They also needed the ability to drill down into specific transactions to research customer inquiries or analyze pricing proposals. "We needed both broad and deep access to sales data," elaborates Byram.

While the Prelude system enabled ISCO to dramatically improve processes and business results in order entry, inventory management, and financials, asking sales reps to use the Prelude ERP application to get the information they needed proved impractical.

"It's time-consuming to log in and out of the ERP application when you only need a certain kind of information" says Byram. "We didn't want to force the sales reps to use their time that way." While a dashboard solution was available from Activant, it did not provide access to all of the types of underlying sales and customer data ISCO sales reps needed.

ISCO also had stringent security requirements that forced it to look for other solutions. "Our primary focus for the initial solution was our sales force. We needed to give each sales rep access to their sales, order, and customer information, without giving them access to info for all the other sales reps." explains Byram.

Usability and a short learning curve were also critical. "With so many reps, each of them critical to our business, we were determined to find a query and reporting solution that was easy to learn and use. Acceptance by the sales force was going to be a key and we couldn't afford to choose anything that would require a lengthy implementation and training phase."

MITS Discover: The Right Blend Of Access, Security, And Speed

In a sense, the search for a business intelligence solution to support the required queries and reports began before ISCO's need was understood. Byram, in a previous position managing Oracle databases, had already invested significant effort evaluating business intelligence and Online Analytical Processing (OLAP) solutions so she fully understood (and embraced) the range of features and capabilities available from leading solutions.

At ISCO, Byram quickly became aware of the MITS Discover business intelligence solution. "Discover was already in use with a large number of Prelude implementations, in something like two-thirds of the Prelude installed base, so we knew it was likely to be a good solution for us."

One of the most positive aspects of implementing MITS Discover has been speed. "It's so easy to learn and use that we usually need less than half of the 30 minutes we set aside to train an end user," reports Byram. Rapid training and user acceptance have also meant a faster overall rollout. Byram adds, "We started slowly, introducing just a single set of sales queries, but things went so well that we very quickly rolled out three additional sets: open orders, open quotes, and accounts receivable."

The results have been gratifying. "Now the sales force can very quickly get all the information they need to do their job, all four parts of the business they need to understand, without having to come into Prelude." Support is also a positive: "We've had very few questions coming back to us from the sales force. They just don't have any problems understanding how to get the information they need. It's a real time saver for IT."

Results: Vision Fulfilled

When asked how ISCO had completed one of the largest and fastest, MITS Discover implementations for Activant Prelude on record, Byram credited the solution itself coupled with MITS support. "Most importantly, we had a vision. We knew how we wanted things to work, we knew the capabilities we needed, and we knew we wanted to take things in steps. Once we saw Discover, we were convinced it would fit that vision nicely." ■

Web Developers Sierra Bravo Moving to New Nerdery (Office)

Sierra Bravo, the largest independent web development firm in the Twin Cities, recently moved into a new office.

"We're moving just blocks away in Bloomington, but to us, our

move looms as both large and symbolic," said Luke Bucklin, president of Sierra Bravo — one of three programmers who founded the company in 2003.

Sierra Bravo now has a staff of over 60, including more than 40 programmers collectively versed in an unstumpable range of programming languages and platforms — an embarrassment of nerdy riches in web-development capabilities.

"With our nerdy-powers in full bloom, we outgrew our space," said Bucklin. "We expect to continue our current annual growth rate of 40-50% for the next few years, and our new space gives us room to keep adding the tech talent our company and clients need."

Sierra Bravo's address is:

9555 James Ave. S., Suite 245

Bloomington, MN 55431

(Phone remains 952-948-1211)

"Actually, who are we kidding here?" added Bucklin. "We'll get all giddy and frolic about our nerdy new space. We'll test each room's acoustics by plugging in Guitar Hero and playing the theme from 'The Jeffersons' like our lives depend on it —because we're moving on up to a brand new Nerdery! Our new office even has showers. We will bid our families a fond farewell until the holidays."

Sierra Bravo's nerd-powered growth is fueled by its knack for resolving business needs through custom software and web solutions that help companies work smarter. Through its Partner Program, Sierra Bravo's web programmers, interface designers, and database engineers add sophisticated online functionality to the creative concepts of design and communi-cation professionals at ad and design agencies.

About Sierra Bravo

Sierra Bravo creates practical, business-minded web solutions that suit clients' real world needs, including e-commerce and content management solutions. As creators of the F1 Overnight Web Site Challenge, Sierra Bravo led 88 volunteer web pros to donate 24 hours of their time and talent to build free web sites for 11 Minnesota non-profits — collectively valued at $200,000. Visit us at www.sierra-bravo.com. ■

Alliance to Energize Fundraising for Higher Education Institutions

Datatel, Inc. announced that it has partnered with WealthEngine.com, a donor qualification and wealth identification provider, to offer colleges and universities an advanced fundraising solution. The agreement allows Datatel Colleague Advancement and WealthEngine's FindWealth software products to share information about potential donors and provide reports via a user-friendly interface. The announcement was made in New York at the annual Council for Advancement Support of Education (CASE) 2008 Summit for Advancement Leaders.

"Institutional development is a key component of any college's or university's Strategic Academic Enterprise," said John Speer, Datatel president and CEO. "Our alliance with WealthEngine gives Datatel Colleague Advancement clients a competitive edge when targeting donors, whose support is critical for student success and institutional sustainability. We look forward to working with our new partner to bring this highly desirable fundraising solution to the higher education market."

"Technology is driving fundraising research towards a more timely and systematic process," said Tony Glowacki, president and CEO of WealthEngine. "Gone are the days when a college or university would wait two or more years to research donors, parents, or alumni because of the expense and time it took to sort through the results. With this partnership, Datatel's clients can determine when — and what parts of — their Colleague Advancement database should be screened for wealth identification, so that it is frequent and fulfilling. Not only are the results immediate, but key data is then imported back into Colleague Advancement software so users can act quickly and, ultimately, be more strategic and successful in their fundraising."

A fully-integrated alumni relations and institutional development solution, Colleague Advancement delivers real-time key performance indicators, executive reporting capabilities, enhanced communications tools, and a host of new features enabled through advanced technology. Colleague Advancement offers hands-on access to key prospect and donor information that enables users to focus on fundraising efforts to achieve their institutions' goals, as well as streamline processes, strengthen relationships, and deliver results.

FindWealth Screening takes donor and prospect lists and data mines them against more than 27 sources, including real estate, philanthropic and campaign gifts, corporate biographies, stock holdings, and other hard asset databases. Data analysis is then applied to the results through statistical formulas and ratings so that each record is segmented for specific fundraising campaigns and development officers know who to target first. The results can be integrated with Colleague Advancement for seamless, up-to-date research.

About Datatel, Inc.

Datatel provides fully-integrated administrative and academic software systems and professional services to colleges and universities, for building Strategic Academic Enterprises in support of student success. Serving five million students at more than 750 institutions throughout North America and overseas, Datatel has exclusively focused on higher education for 29 years. Datatel is headquartered in Fairfax, Va. For more information, visit www.datatel.com.

About WealthEngine

WealthEngine offers Web- and software- based data mining technology to nonprofits and financial services organizations. More than 1,600 clients use WealthEngine for comprehensive research on individuals, companies, and foundations. Research from more than 27 different databases is bundled into handy profiles containing hard asset, charitable and campaign giving, real-estate, corporate and other biographical data. Headquartered in Bethesda, Md., WealthEngine offers products in both the United States and the United Kingdom. For more information, visit www.wealthengine.com. ■

SJ+ Integrates Entrinsik's Informer Reporting Software with PRC Solution

SJ+ Systems Associates, the leading supplier of software development lifecycle and IT governance software for MultiValue users, has announced beta availability of its PRC solution with Entrinsik's Informer web reporting utility, the industry's most popular interactive reporting tool. For customers using IBM UniData and UniVerse (U2) and other MultiValue databases,

PRC provides a framework for flexible and extensible configuration management.

With the addition of the web-based Informer tool, PRC customers will experience a fully-integrated, intuitive platform for reporting, enabling them to easily access data directly from PRC and view the data in HTML, Excel, or PDF formats and e-mail reports. They will also be given the option to extend Informer reporting capability across the rest of their databases.

"We are delighted to have SJ+ as a partner, and that PRC users will now benefit from this exposure to Informer," said Entrinsik President, Doug Leupen. "We look forward to working with them to expand their customers' reporting capabilities."

Susan Joslyn, SJ+ President and author of PRC, says, "Some of our newer customers consider

auditor's reports as the primary deliverable of their PRC implementation, while others have simply not been taking advantage of the wealth of information and metrics being collected. In either case, the addition of Informer can help them easily create attractive, functional, and useful reports that go beyond mere compliance. We were impressed with the beautiful job Entrinsik did with the tool, and are pleased to offer it to our customers."

About SJ+ Systems Associates

SJ+ Systems Associates develops and supports PRC, a mature, complete software configuration management tool for U2 and MultiValue applications, including SB+. SJ+ has the experience and education in IT governance, audit, software quality, and best practices to help customers implement a complete

lifecycle management strategy that is compliant and productive. For more information, visit http://sjplus.com.

About Entrinsik, Inc.

Founded in 1984 and in the MultiValue market since, Entrinsik, Inc. has become an industry leader in delivering reporting systems for the IBM U2 family of database products. With Informer, thousands of end-users using MultiValue databases leverage the Web to securely share information anytime, anywhere. With its intuitive, easy-to-use interface, Informer provides end-users, IT staff, and executives an operational business intelligence environment by delivering up-to-the-minute information impacting key business processes. Entrinsik provides leading edge technology and unsurpassed service to their customers. For more information visit Entrinsik at www.entrinsik.com. ∎

Clif Notes

the Moon. Whichever it is, the other IT techs are giving accounting a wide berth for the time being.

Even though I've been asked the question about every three weeks for several years, I've never come up with a good answer to, "Why are the people in IT so incapable of saying 'I don't know,' and why do they feel they have to just make up answers to avoid getting caught?" I suspect we have all done it at some point in our careers. I'm pretty sure I went through my period of being an arrogant little twit. One of the nice things about being in your 50s is that you tend to forget what you were like in your 20s. Assuming I am right, I've tried to figure out what, other than age, made a difference. Becoming self-employed probably had a lot to do with it. If you tick off a lot of lowly users and get banned from the site, you can starve rather quickly. Getting paid or not getting paid is a pretty effective way of driving home the point of just who your "customers" really are.

What if your IT department docked $100 from the paycheck of the IT techs for each complaint from a different user that pay period. Or what if they had a "three strikes" rule? The third time you are caught making up a phony answer to a user just so you can avoid saying "I don't know, yet," you're fired. Or to put a more positive spin on it, how about management polling the users every quarter as to which IT tech they thought tried to help them the most? That tech gets a bonus.

I'm not sure I have a really good answer. I am sure, however, that this snotty war we, as a group, seem to have going on with the users has got to stop. After all, these people are the ONLY reason you and I have jobs.

Of course, I suppose I could contract my daughter and her friends to come out for a couple of months and "touch up" your IT department.

Be afraid. Be very afraid. **IS**

TOP 5 CUSTOMER SERVICE MISTAKES IT PROFESSIONALS MAKE

and project management work, I run into this one often. While I no longer program or install hardware, I have a strong understanding of technology, business practices, and programming logic. It doesn't help me to tell me it can't be done, because I'm going to be looking for a way to accomplish it. I need more information, more of an explanation, of why something is happening than the average customer. If you have customers like me, it is helpful to give a little more.

I wish I could give you exact times of when to give more or less of an explanation. But there isn't really that simple. It is an art, and depends greatly on who your customer is. Here is where it pays to build some sort of (positive) relationship with your customers and learn their preferences.

Now, our next mistake is probably the fuzziest of all of these five, at least as far as how to best deal with it. But I have to mention it.

Don't Have Time to Deal with It

Mistake number five: Telling your customer "I don't have time to deal with it."

We all deal with overloaded workloads at one time or another. It's a good predicament usually. But telling our customers that we "don't have time to deal with" their request is the same as saying "NO" or shutting the door in their face.

Do you want the phone company or the cable company to tell you that you can't have service because they don't have time to set it up for you? No. You don't want your customer in the same predicament. We need to find an alternative.

Yes, you have too much work right now, but you can say "Listen I'm really swamped right now, can I call you tomorrow about this?" Most of the time

your customer is going to agree. But if they have a hot issue, then you need to either personally hand them off to someone who can help them, or find a way to reprioritize what you are working on.

So what do you do when you have multiple high priority items competing with each other? You communicate, with your manager, with your customers. Some things are more critical than other critical items, and some customers are higher profile than others. Sometimes internal customers (project managers, managers, executives) will work out the priority issue among themselves and solve your dilemma. But you never tell your customer they are out of luck because you can't deal with it, not without giving them another solution. That's a really good way to be forced to make a career change.

Continuous Improvement

So did you find anything that you need to work on? I know I did. All of this boils down to continually improving on our communication and interpersonal skills. As most of us in IT are introverts, building these skills is usually outside our comfort zone. Believe me, I understand. I've spent the last 15 years focused on building these skills, and I am still working on them. But I also know that working on one area at a time and improving these skills even a little can make big difference in our careers and businesses. **IS**

SHANNON STOLTZ is an IT communication consultant and freelance writer. With over 15 years in the IT industry, Shannon has cultivated a customer-centric focus through a variety of roles including training, project management, software installations, technical writing, governance, business analysis, public relations, and copywriting. Shannon can be contacted at shannon@shannonstoltz.com.

PREPARING FOR BUSINESS INTELLIGENCE

Continued from page 18

The order of the steps is important — it is common (and risky) for IT staff to jump straight to choosing a source file because they make that choice instinctively as a way of thinking about the project; while this is frequently correct, the downside of choosing the wrong file can be substantial.

The HS.SALES account will provide examples for the four steps.

Formulate a Vision Statement

Before any project (not just a BI project) it makes sense to set out a basic vision statement outlining the business process being modeled and the shared vision of the scope and capability of the finished project. This helps in two ways: first, without a clear idea of the scope of the project, it is difficult to know when the project is finished. Second, developers don't need to go back to the client for every micro decision if the vision statement is clear enough. If the vision statement includes the phrase "fast and cheap" little time will be spent choosing the colors for the staff's Aeron chairs. If decisions are outside or contrary to the vision statement this is a red flag to rethink or reset expectations — early.

A vision statement should be crafted in consultation with the intended audience. Understand their goals and how this project will support those goals. It is sometimes helpful to collect questions to be answered by this project. A key question can be used to validate the design: "which products receive the highest discounts?"

Reach agreement with all interested parties. This vision statement will be used for the rest of the article: "Provide upper management with information on sales and discounts by customer and product." This simple statement includes the audience, the business process, and specific requirements.

Choosing the Transaction Granularity

Transaction granularity refers to the atomic level of the data. Each line in a flat file is at this level of granularity. Baseball standings have team-level granularity. Line scores have inning granularity. The example vision statement looks for "product sales." This leads directly to a "line item" granularity — products or SKUs are on invoice lines. Had the vision statement been restricted to "customer profitability," then "invoice" or "customer" granularity would be sufficient. Three factors come into play:

1. A more summarized granularity will create a smaller, less resource intensive solution;

2. A more detailed granularity will provide more flexibility; and

OPTIMIZING UNIDATA BASIC VIA CODE PROFILING, PART 3

```
ANS = OCONV(VALUE,MR20') ;* Show 2 decimals, scale 0 decimals
```

Fig. 13 Alternative code with less decimal movement

back to the left two places and show the value with two decimals."

As an alternative, why not use the code in figure 13, which achieves the same objective without so much decimal movement or CPU usage. This is, of course, just one simple example, and yes, the amount of CPU saved is small, but when it's your CPU and your users, every CPU cycle you save may make just the difference they need. **IS**

KEVIN KING is the President and Chief Technologist with Precision Solutions, Inc., a leader in technology solutions, support, and training. He is also the author of SB+ Solutions, an enthusiastic private pilot, and Christian guitarist and producer... as time allows.

Note

Windows proponents should not use these tests as a representative comparison between the different platforms, as the configuration and load between the Windows and AIX servers used in this test are dramatically different. What should be gleaned, however, is that Windows may be more susceptible to being pegged by a single rogue program, and yet due to the way that Windows schedules the CPU, the statistics that should tell us there's a problem are effectively defeated.

PREPARING FOR BUSINESS INTELLIGENCE

Continued from page 26

3. Matching the granularity of existing models can potentially provide extended capabilities and help validate the accuracy of the solution.

Granularity and Multivalues

If the granularity of the facts is at a multivalued level, all fact fields must be associated multivalues and any multivalued dimension fields must be associated. Non-multivalued fields must be repeated to flatten the data. For example, if five multivalued line items are associated with five sets of SKUs, quantities, and prices, a single shipping charge cannot be included as a fact. A shipping method for the entire order could be included as a dimension field, but it will need to appear five times in the flattened data. Artificial facts like line counts would appear as five associated multivalues, in this case, each with a value of 1.

Choosing a Source File

With an understanding of the transaction granularity, the primary source file choice is usually straight forward. There will be supporting files, but any SELECT and extraction will be based on this primary source file. An "invoice line item" granularity will typically only yield a few choices for a typical MultiValue system. The history file may contain all invoices for all time — but only contain summary data. An open orders file may have the appropriate granularity, but not the required history. Another file might contain both current and historical transactions, but be highly volatile and difficult to validate.

If there are multiple candidate files, consider choosing one that is already used in the reporting solution or one that has an existing body of reports. This will make it easier to audit the solution.

The ideal is a file that best meets the vision statement while supplying the requisite transaction granularity. If there is no appropriate source file, do not proceed. Revisit the vision statement with the project stakeholders to see if a compromise with existing data is possible, or revise your application to start capturing the required data going forward. There is little point in trying to build an analytical reporting solution without the data to do the job.

The example vision statement leads to "line item" granularity. This level of granularity is actually stored in a sales history embedded in each CUSTOMER record — and CUSTOMER will be the primary source file. Can there be multiple source files? Yes, but usually only one primary source file matches the required granularity. It may be logically or physically distributed, in which case the files need to be merged as part of the data extraction.

Choosing the Dimensions and Facts

Stepping through the fields in the CUSTOMER record, the vision statement will guide the treatment of each field and determine what part it plays in the dimensional model.

Customer ID: 1

Include in Customer Dimension. This allows upper management to locate transactions for a particular customer. In a star schema, this will be the key to the customer dimension.

Contact Name: Mr. Samuel Smith

Do not include. Upper management will not be interested in this level of detail.

Company Name: Better Beer, Inc.

Include in Customer Dimension. This allows upper management to locate transactions for a particular customer.

Street Address: 10 Commercial St.

Do not include. Upper management will not be interested in this level of detail.

City: Concord

Include in Customer Dimension. This allows upper management to look for geographic patterns.

State: NH

Include in Customer Dimension. This allows upper management to look for geographic patterns.

State name: New Hampshire

Include in Customer Dimension. This allows upper management to look for patterns. Do not include this in a separate "State" table. This is what is known as snow-flaking dimensions and leads to exceptionally slow query speeds.

Zip: 02131

Include in Customer Dimension. This allows upper management to look for geographic patterns.

Telephone: (603)555-3212

Do not include. Upper management will not be interested in this level of detail.

Product: M2000

Include in Product Dimension. This allows upper management to locate transactions for a particular product. In a star schema, this would be the key to the product dimension.

Description: Moderate duty, monochrome copier

Include in Product Dimension. This allows upper management to look for patterns.

Serial#: 501278

Do not include. Upper management will not be interested in this level of detail.

Date Purchased: 01/07/91

This is the time dimension. Every transaction-based dimensional model will have a time dimension.

Date paid: 01/28/91

Do not include. The amount of time it took customers to pay has no bearing

on the vision statement. The fact that they paid might.

List price: $4,490

Include as a fact.

Price: $4,200

Include as a fact. Combined with List price, yields discount percentage at any level of aggregation.

Discount: 6.5

Do not include as a fact. Discount percentage is a semi-additive measure. These percentages cannot be added, but can be calculated from the sums of list and price. Individual order discount might be used to drive a discount cohort dimension where ranges of discounts are grouped for analysis, e.g. 0%, under 5%, under 10%, over 10%.

The same dimensional model yields a "flat file" representation (fig. 3) or a "star schema" (fig. 4). Before moving on to extract and transfer the data from MultiValue, this is a good spot to "sanity check" the dimensional model against the vision statement and primary question. This "sign-off on the blueprints" makes an excellent milestone — and can potentially save a tremendous amount of effort. In this case, the goal was to calculate discounts by product. With this model, this is possible at any level of detail — as long as our reporting tool allows us to do calculations on intermediate totals in columns

There are other "non-transactional" dimensional models, and additional complexities caused by stretching a transactional model that captures the truth at an instant to fit a reporting model that accurately shows a flow over time. These will be covered in future articles. **IS**

ROSS MORRISSEY helps MultiValue shops with Business Intelligence challenges. If you're stuck, can't get started, or want to do more; contact ross@tantiva.com.

```
Flat File Solution:
Customer ID.... 1
Contact Name... Mr. Samuel Smith
Company Name... Better Beer, Inc.
City........... Concord
State.......... NH
State name..... New Hampshire
Zip............ 02131
Product........ M2000
Description.... Moderate duty, monochrome copier
Date Purchased. 01/07/91
List price..... $4,490
Price.......... $4,200
```

Fig. 3 Flat file solution
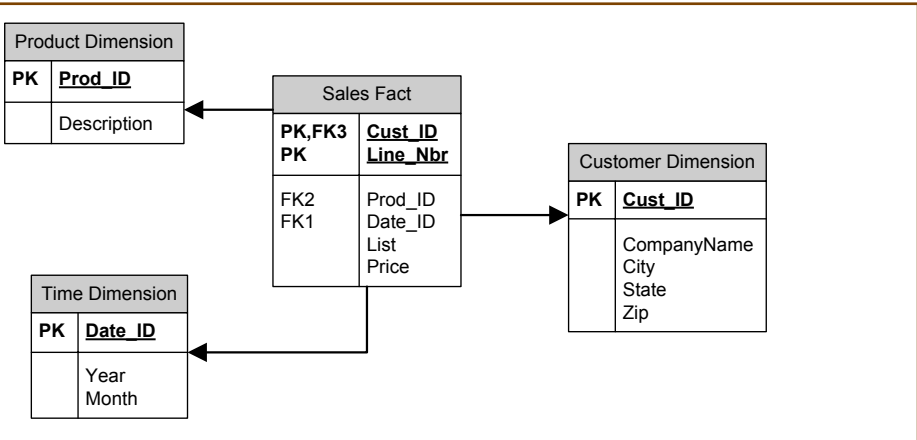
Fig. 4 Star Schema solution

BY MIKE RUANE, REVELATION SOFTWARE

Get the Picture?

Using Graphics in OpenInsight

One of the great parts of using Graphical User Interfaces, or GUIs, is the fact that they are *graphical*. You can use graphics on the screens and reports, instead of trying to make do with text. The possibilities are obvious and exciting: pictures of items in a catalog, employee picture IDs, maps of work sites, and more.

Using images in OpenInsight is easy, and will be even easier in future releases. In this article we'll instruct you on how to make good use of graphical images in OpenInsight applications. The content provided here applies to OpenInsight versions 8.0.7 and above, although nearly everything works in prior versions.

What Graphics Formats Does OpenInsight Support?

As would be expected, OpenInsight supports nearly every type of graphical format for images that is out there. But first, why are there so many different graphical formats? Why are there so many flavors of MultiValue? The answers are probably the same: They are designed for different purposes, and each particular vendor thinks theirs is the best.

Most graphics files are identified by their filename extension. JPG or JPEG files, from the Joint Photographic Experts Group, are well known, used widely, and has its own compression. BMP files, or bitmap files, were developed for OS/2 and Windows and support up to 24 bit images. GIF, or Graphics Interchange Format, is efficient but only supports 256 colors. PNG, for Portable Network Graphics, supports from 1 to 48 bit graphics. There are more, some less popular than others, and some older than others.

As graphics formats evolve and become more widely supported, Revelation Software expands their support for them. However, once a format has been incorporated into OpenInsight, support for it will likely never be dropped.

Graphics in Forms

So, now that we've covered the different types of images that OpenInsight supports, let's discuss how they can be used in an OpenInsight form. As has been described in numerous articles in International Spectrum, an OpenInsight window is comprised of controls placed on a form. There are many supported control types such as Edit lines, Edit Tables (for multivalued fields), Combo Boxes, and others, including Bitmap controls.

Although they are named Bitmap controls, these controls are used for all of the supported image formats. Adding a Bitmap control to an OpenInsight form is simply a process of choosing the Bitmap control from the Form Designer toolbar, as seen in figure 1, and placing it on the form.

Once the bitmap control has been added to the form, its default property panel opens, as seen in figure 2. The image with the phrase 'Bitmap' and four color rainbow is the default bitmap image loaded by OpenInsight.
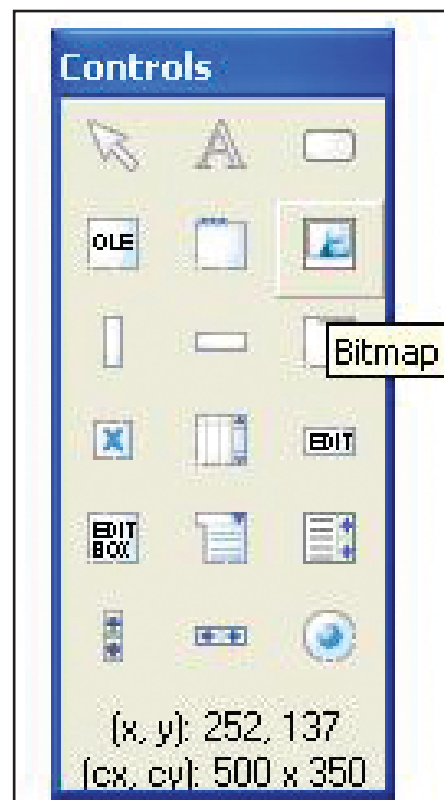
Fig. 1 The OpenInsight Form Designer Toolbar

As can be seen, Bitmap controls do not have a lot of properties. They have a name and a number of styles relating to whether they are visible, should appear on all pages, and should automatically adjust their size as the form does, or should they be anchored. There are also three mutually exclusive styles regarding how the image should be resized within the control itself.

Graphics appear in all sorts of sizes but mostly two: square (think images on buttons or icons) and rectangular (think of pictures, 8x10 or 5x7 inches). The chances that the size and shape of the Bitmap control on the screen will match the size and shape of the graphic that should be displayed are usually slim to none. Therefore, OpenInsight offers three ways of reconciling the difference: clip the image, stretch the image, or scale the image.

Clipping will take the original image starting at the upper left corner, and show enough of the image to fill the control on the screen. Stretching will distort the image so that it will fit the control completely. Scaling will resize the image so that it scales proportionally to the size of the control. Figures 3, 4, and 5 show the OpenInsight logon splash screen clipped, stretched, and scaled, respectively.

The observant reader will have noticed that this author did not describe how the OpenInsight logon image was assigned to the bitmap control. This can be done in two ways: at design time in the form designer, which we will refer to as a static assignment, and during runtime, based upon user actions or data, which we will refer to as dynamic assignment.

Assigning Images in Forms: Static and Dynamic

Images can be assigned to a Bitmap control either at design time or during runtime. Static images, assigned during design time, could include a corporate logo, an instruction sheet, or something
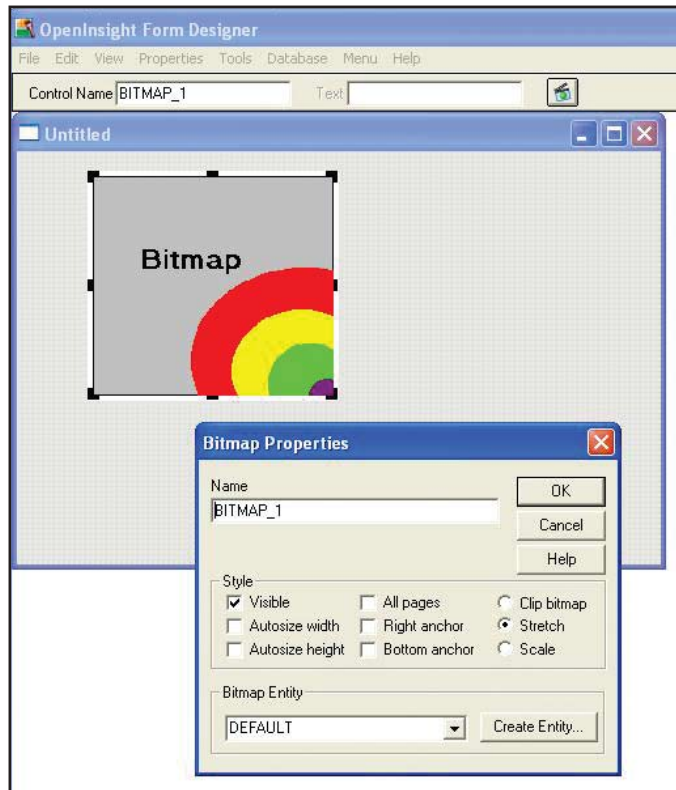
Fig. 2 The Bitmap Control property panel

Fig. 3 A clipped Image

Fig. 4 A stretched Image

Fig. 5 A scaled image
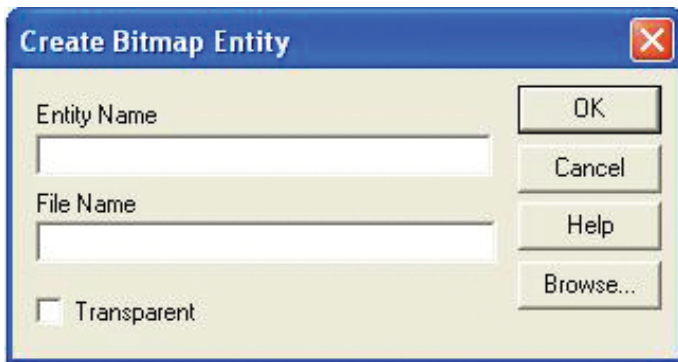
Fig. 6 The Create Bitmap Entity window
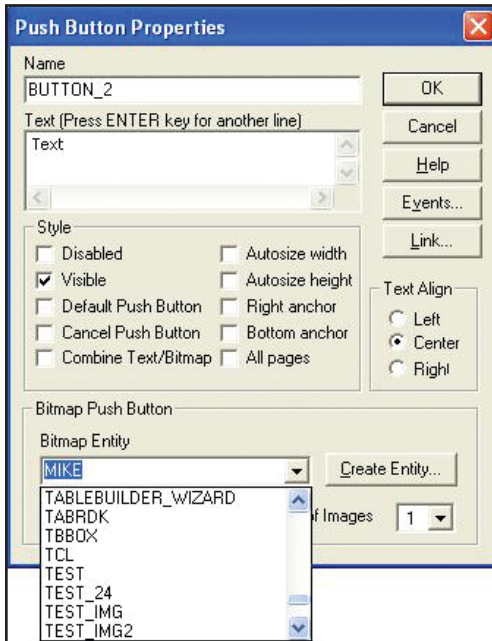
Fig. 7 A list of defined Image Entities
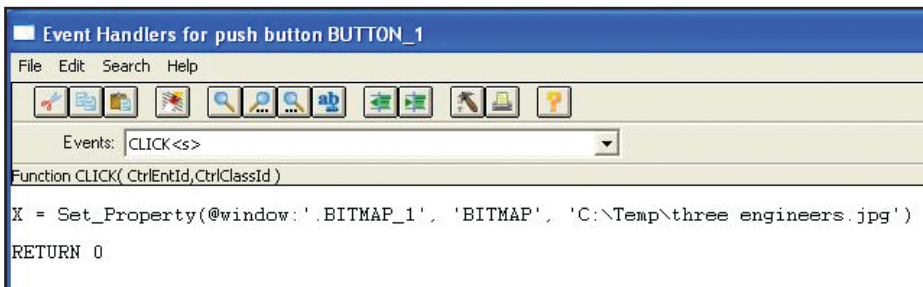
Fig. 8 Code to change an image dynamically

Fig. 9 The screen with the new image

GET THE PICTURE? USING GRAPHICS IN OPENINSIGHT

of the like. Connecting an image to a Bitmap control at design time requires that the image be assigned an Entity Id within OpenInsight.

Creating an image entity is a simple process. From the OpenInsight bitmap property panel, clicking on the Create Entity button will cause the Create Bitmap Entity window to appear as seen in figure 6.

By assigning the entity a name and then choosing the image file, a new entity is created. By default, it is automatically added to the OpenInsight's Repository and is assigned as the Bitmap Entity for the control. Clicking on the drop down arrow will list all the Image Entities currently defined in the system. Figure 7 shows an example of this list. Saving the form and test running it will cause the control to display whatever image was selected.

Assigning images dynamically requires a few lines of basic code in most cases, but the idea is extremely simple. By setting a Bitmap control's BITMAP property to the name and location of an image file, the image control displays the image. I want to repeat that even though the control is called a Bitmap control, and the property being set is the BITMAP property, the control is used for all of the supported image file types.

The images can be displayed either based upon system actions or user actions. An example of a system action would be that after reading an employee's record into a screen, their ID picture is displayed. An example of a user action might be scrolling through an inventory.

The code required for this process is very simple. The image path and filename is assigned to the control, and the control displays the image — either clipped, stretched, or scaled — as defined during design time. Figures 8 and 9 show a code example and the resulting screen.
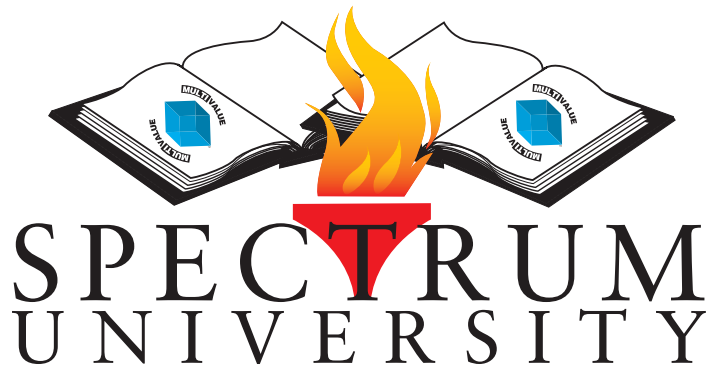
Upcoming Webinar Schedule

TIME AND DATE	TITLE
Tuesday **September 23, 2008** **2:00 pm EST, 11:00AM PST**	**Back-to-Basics: MultiValue Dictionaries** If you use MultiValue Retrieval in any aspect of your job (end-user or developer) you need to understand what is happening in your file dictionaries. During this 6 week session you will learn how to read and write dictionaries, understand the various dictionary structures and how they process. We will cover what each line of the dictionary does, how to convert data and create calculated dictionaries.
Thursday **September 25 , 2008** **2:00 pm EST, 11:00am PST**	**D3 Administration** The MultiValue Administration Series is aimed at those who are involved with day-to-day administration of the MultiValue database environment. Each course will be 4 weeks long and will include the following key topics: file maintenance (creation, sizing and resizing), managing users, managing printers, managing foreground and background processes, backup, locking, working with persistent and non-persistent connections.
Thursday **October 23, 2008** **2:00pm EST, 11:00am PST**	**MultiValue BASIC Debugging** Debugging can turn from a science to an art really quick in today's business systems. With the demands on new development, debugging and maintaining existing programs require you to spend time understanding the best way to debug problems. This class will teach you effective debugging solutions using both realtime and passive debugging techniques. This is a class that is a must for all developers.
Tuesday **November 4, 2008** **2:00pm EST, 11:00am PST**	**Integrating Email** Email is part of the everyday business environment. Everyone recieves it and send it, but does your business effectively use it? This class will focus on how to send and receive and use email within your business application. It will show you how to access incoming email, and then how to use that within your system. It will also show you how to create and send emails from your business system without having to retype emails or cut and paste information into your email clients.
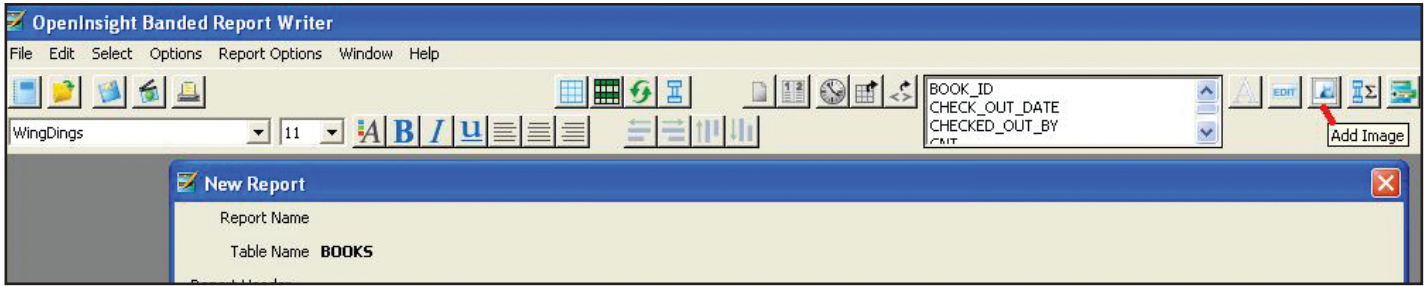
Fig. 10 The Add Image button

> *" Whether on a screen or in a report, a picture is often worth a thousand words. "*
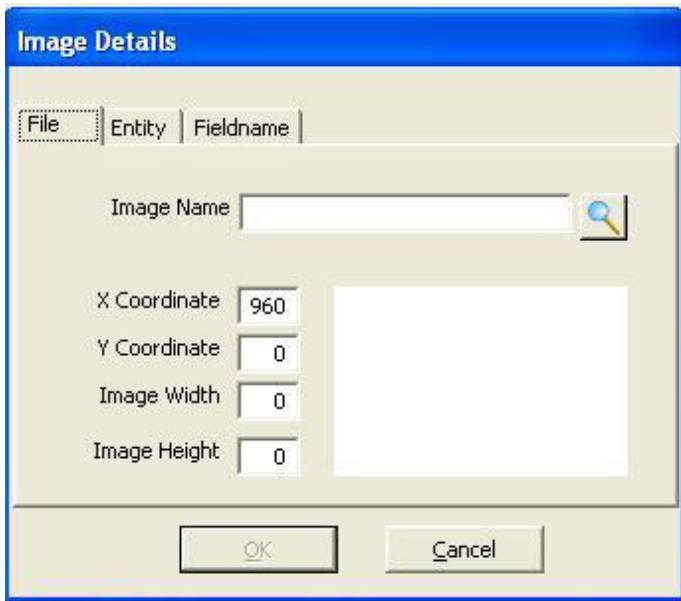
Fig. 11 The Image Details window

Graphics in Reports and Documents

Given the ease of adding images and graphics to screens and forms within OpenInsight, it should also to be easy to add images and graphics to reports and documents, correct? And it is.

In OpenInsight's Banded Report Writer (BRW), an image can be added to a report by clicking on the Add Image button as seen in figure 10. This calls up the Image details window, as seen in figure 11. In the banded report writer, an image can be static or dynamic, without writing any code. At design time, the user can choose either the name of a file, an entity already defined in the application, or the contents of a field. So, if an Inventory File record has a field named 'ITEM_IMAGE', and it is set to the file f:\inventory_pics\item_9999.jpg, when the report is run, that image will appear in the same detail area as the data from the record. Watch for a future article for more detail.

In Basic+, OpenInsight's Basic-based programming language, our OpenInsight Printer Interface, known as OIPI, supports the printing of all image formats supported by OpenInsight. In general, OIPI is instructed to print a graphic with a message or parameter of 'BMP'. Information about sizing and scaling follows, as well as an indicator as to whether or not the image should appear on every page. Figure 12 shows a simple program that prints a sample graphic.

OIPI also supports line draw characters, rectangles, ellipses, line styles, and much more. User can create forms and areas for filling in without having to use preprinted paper. For more detail, look at the two example OIPI programs supplied with
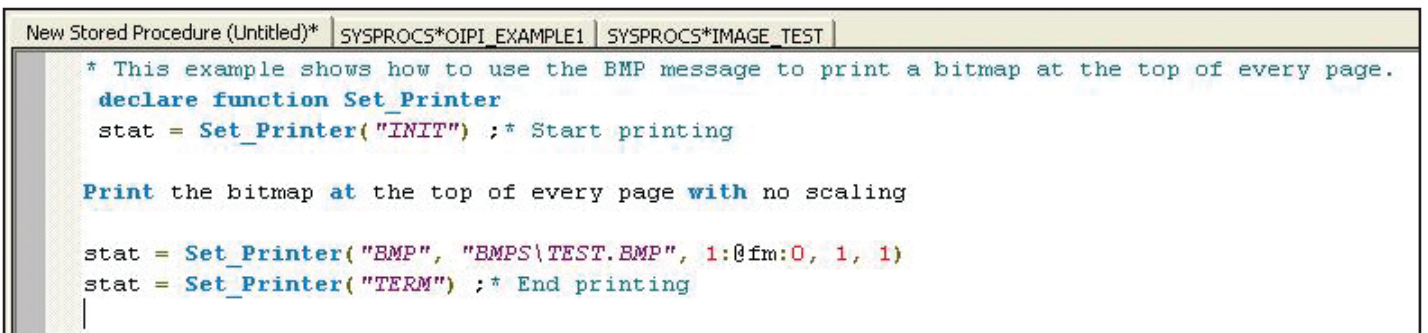
Fig. 12 A simple OIPI-based program showing how to print a graphic

OpenInsight, named OIPI_EXAMPLE1 and OIPI_EXAMPLE2.

What About Moving Pictures?

So far we've only discussed still images, which work well for many applications and are a great improvement over character-based applications. However, there are certain times when having moving pictures (i.e., videos of some sort) is helpful. Does OpenInsight support them?

Not directly, but because OpenInsight's Form Designer supports the use of OLE controls, a video playback control can be embedded in OpenInsight. At this point playing videos becomes a simple exercise. In fact, Window's Media Player can be launched from OpenInsight or embedded within a screen, so all your bases are probably covered. Read our OLE article in the July/August 2004 issue of International Spectrum.

Security

In most of the example described earlier in this article, the image files are located in directories in the native formats, such as JPG or BMP files. There are cases though where this is not really acceptable — medical records, copyrighted images, signatures, and so on. One solution is to protect the directories where these files are located, but that can sometimes become cumbersome. Another solution is to store the images within OpenInsight's Linear Hash filing system.

The process would go something like this:

1. Original Image is OSREAD into memory.

2. The image is ICONV'd with a type of VB, for Variable Binary.

3. The image is saved in a record in an OpenInsight table.

4. When the image needs to be displayed, the record is read.

5. The read image is Oconv'd with a type of VB.

6. The image is sent to a Bitmap control, setting the control's IMAGE property to the Oconv'd value.

This makes for a very secure image storage paradigm. If anyone were to somehow crack the system security and get a hold of the data files, all that would found are a string of ones and zeroes. As added precaution, a developer could add their own encryption routine to the mix. A neat and elegant solution to what can be a sticky problem.

Conclusion

So, do you get the picture? With OpenInsight, it's easy to do so. Whether on a screen or in a report, a picture is often worth a thousand words. With OpenInsight, you can speak volumes. **IS**

What Good is IT Governance, Anyway?

BY SUSAN JOSLYN

For the last several years the talk in IT has been "all about compliance." Finally things have returned to some semblance of (over-audited) normalcy and now the talk is "all about value." In the natural way that things progress, what began as a study of the value of all this compliance evolved into an examination of the value of IT governance in general and that naturally progressed to an examination of the value of IT itself. For those of us in IT the only difference is that it looks like folks on the other side of our partition walls might be interested in the answers this time.

Many of us in the technology side of the world first encountered IT Governance in the context of audit guidelines due to new compliance requirements. Sarbanes-Oxley and changes to HIPAA as well as other governing bodies have brought the maligned and obscure practice of IT auditing to the (no less maligned) forefront. But the context has uniformly been about compliance. It's been about ticking the box and following the rules.

Interestingly, in a parallel universe, Gartner and Standish have reported that 66% of IT projects fail on some level. The first reason cited is poor project management. The second reason? You guessed it — poor IT governance. We know that the catchall of project management has probably always included IT governance, but here it is getting its own column in the press. Now it is being suggested that IT governance be budgeted into every project. To which we respond with a resounding wail,

"Yet another weekly meeting!?" Relax, we don't need another meeting. Budgeting IT governance into projects can start out (and stay) as simple as allowing the project time to go through the proper life-cycle, keeping the metrics we're meant to keep, *and then evaluating the project against those metrics*.

It is human nature to not look back at projects. We succeeded? Yeah! Move on! We failed — bummer, let's not think about it. But that approach misses the whole point of doing Project X in the first place: solving a business problem. After all, not a single IT project would ever happen if there weren't a business-driven need. If we don't look back to see if we met the need, and how effectively, then we can't really call it a success. In order to measure success we need measurable benchmarks and then we must actually measure our activities against those benchmarks.

While there are no one-size-fits-all best practices per se to finding business

value like ITIL for service management or CoBIT, human nature probably plays a bigger role in the benefits realization game than any other factor. Research has found that the few companies that do consistently look back at projects to see if they actually delivered any value stick to the basics. They look at the problems they want solve, agree on the metrics that will indicate a solution has worked, and then go back to see if said solution did the trick. That doesn't sound that hard, does it?

Taking the time to look back helps us see value which we can present to the business. We want them to see that we are providing value — and we need to show them that in a way that they can appreciate. Why does that matter? Because what we really want — besides credit now and again — is to be included in the planning and decision making. We want to be respected for our past contributions and counted on for our future contributions. Also, each time our contribution brings value we want that value to be visible.

Not taking these steps to measure IT value so that we can recognize its benefits and to ensure our seat in the decision-making conversations is costing many of us. According to a study by Diamond Management & Technology Consultants Inc., of Chicago, 87% of business leaders say they believe that IT is critical to their companies' strategic success.

However the same study finds that few businesses have yet positioned IT in a way that allows it to achieve this. Only

INTERNATIONAL Spectrum®

MultiValue Business Developers Conference and Exhibition

The Westin Westminster | Denver, Colorado

SAVE THE DATE!

March 23 - 26, 2009

For more information visit www.intl-spectrum.com

WHAT GOOD IS IT GOVERNANCE, ANYWAY?

Continued from page 36

> "
> *Budgeting IT governance into projects can start out (and stay) as simple as allowing the project time to go through the proper life-cycle, keeping the metrics we're meant to keep, and then evaluating the project against those metrics.*
> "

33% of business leaders say IT is very involved in developing their company's strategy, and 30% say the business executive in charge of strategy works closely with IT. That has an impact on performance: 76% of companies say they have had to abandon a tech project, including 29% that say they have abandoned more than 10% of these projects.

At the same time, some companies are reaping benefits from taking steps to measure IT value. At the pharmacy chain, Walgreens, the CIO has been included on the management team since the 1990s. This has allowed Walgreens to connect all their pharmacies with a single computer system that is updated regularly based on feedback from its pharmacists and employees.

What can we do to realize value from IT and specifically from the IT governance that we implement?

David Sward, Senior User Experience Researcher at Intel Corporation, blogs about Intel's business value program (intl-spectrum.com/s1004). He says that the key to explaining and measuring the business value IT solutions deliver is to use a "common language" that is defined *not by IT*, but by the customer. To do this, you need to know what your end users, customers, and company as a whole deem valuable. "For example," Sward explains, "employees responsible

for product inventory won't think of IT solutions in terms of server uptime, database optimization, etc. They want to know specifically how the IT solution is going to allow them to better manage inventory and to do their job better and more easily."

On his blog, Sward explains that creating business value "dials" provide common reference and clarity on how IT solutions deliver value. Using this common reference, there is a framework for understanding and quantifying IT's business value, which then allows IT groups to "more effectively communicate the value they deliver." From here it is easier to build business cases for IT solutions.

A business case (according to Wikipedia) is the concept of having a non-technical justification for a project or task so that any time resources such as money or effort are consumed, it is in support of the business. An example would be a software upgrade that would improve

system performance, but the business case is that improved performance improves customer satisfaction. (Of course "customer satisfaction" is broad and generic, so a more specific value is better — more on this in a moment.)

Central to both the value dials and the business case is the creation of *operational definitions*. David Sward explains, "this is the step that translates the concept of business value into some type of measurement. Operational definitions are necessary to measure business value. Without them, everyone may be measuring but it is highly unlikely they are measuring the same thing. Operational definitions make measurement independent of a person or group, repeatable by others, and standardize results."

When creating operational definitions, the more specific and concise, the better. As in the example above, *customer satisfaction* is broad and can mean many things to many people. But *per-*

centage of repeat customers is measurable and specific, and directly supports customer satisfaction. Operational definitions bring specificity to these broad concepts.

Now with standardized measurements and a common language, a business case can be developed that is understandable and useful to business and to IT. Business cases can range from comprehensive and highly structured, as required by formal project management methodologies, to informal and brief, such as the example above.

Information included in a formal business case could be the background of the project, the expected business benefits, the options considered (with reasons for rejecting or carrying forward each option) along with the expected costs of the project, a gap analysis, and the expected risks. Consideration, should also be given to the option of doing nothing including the costs and risks of inactivity. From this informa-

tion, the justification for the project is derived.

When this information is delivered using agreed-upon and understood measurements and language, real decisions can be better made — and the same measurements and common language can be used to evaluate the success or failure of the project, after the fact. This becomes even better *real* information for the next decision that the business and IT may make together. And making those decisions together is the true goal of this whole endeavor. **IS**

SUSAN JOSYLN is the author of PRC, a complete, integrated Software Development Life-cycle Management tool for MultiValue-based IT. She has worked with MultiValue software since the beginning (hers and MultiValue's) and has specialized in quality, compliance, and life-cycle productivity issues since the early 1990s.

PHANTOMS DO MORE THAN GO BUMP IN THE NIGHT

on the PC. The PC was on the network with the UniVerse system and the phantom polled that file so that when POs were received, the basic program, running as a background process, created an order and sent a notification to the appropriate customer service rep. This saved the EDI person from 1) printing each PO, 2) looking up the customer to determine what service rep it belonged to, and 3) getting up from her desk to deliver the order.

- Another customer would have our customer service rep get on their web site and download their orders as HTML documents. We had one person in the office get on their site a couple times a day and download all orders. The orders were for various inside salespersons so the phantom process could notify each person that an order had been received. Not only would this assure that the notification of a new order was passed on, it didn't require each rep to remember to check the Web throughout the day.

- I had a customer where I installed VSI-fax to automatically fax POs to their vendors at the end of purchase order entry. Though each desktop could have had the client software to view the faxes and e-mails that had been sent, it wasn't acceptable to expect the buyers to check the

whole list periodically to see if any of their faxes failed. The solution required both the Windows scheduler and a background process. First the phantom would write a batch file (.bat) that issued VSI-fax command line commands to request the log of all faxes scheduled since the last request. (The MultiValue system stored the last time it requested that log.) The VSI-fax commands would redirect the output to another Windows text file on the server. The Windows scheduler was set to run that batch file periodically. Then the phantom would read the log that was created when Windows executed the .bat file and search for all failed fax attempts (after three tries). If a fax failed, the phantom would send an e-mail back through VSI-fax to the buyer who sent that PO.

Today's phantoms need not be confined to the night. They bring us into contact with a world beyond MultiValue and make the interface easy and friendly. I hope this article has given you some ideas about when you might find an encounter with a phantom an enjoyable experience — happy hunting! **IS**

CANDI HART has been an independent consultant in Southern CA since 1980. She was known as Candi Piech when she served as president of CDBMA. She may be contacted at candihart@sbcglobal.net

BUSINESS TECH: THE ART AND SCIENCE OF TRAINING (PART 1)

- If you are bad at telling jokes, don't.

- If you have to cut content to fit the jokes in, don't.

- A great deal of humor is 'victim' humor. It speaks negatively about a particular race, creed, color, category of employee, etc. This is a remarkably poor choice, especially if the presentation is being recorded for a wider audience. Not only will you look bad, but everyone who laughs with you will look bad as well.

- Make sure the 'bit' has a point which is relevant to the topic.

When I do use humor, I try to make it inclusive, gentle, and observational. Many of my presentations opt out of humor entirely. Rule 4: *If laughter is the best medicine, then use it in controlled doses only.*

The Sinker

Unlike some trainers I know, I always make a point of attending training myself. I never assume that my ability to present as an expert in one topic reduces my need to be a student on another topic. I have always been comfortable with words like "I don't know" and "I need to learn more about that." So, I've seen a lot of other trainers at work. Many of them know how to set the hook, know when to run a line or not, and unfortunately, don't know how to close.

Just like in journalism, where an article shouldn't just stop abruptly, a presentation shouldn't simply run out of words. It needs a wrap up, a summation of sorts. The good news is that there are several good ways to close. Here are just a few, and you might want to use more than one:

- A recap allows you to give the audience a sense of scope: "In this last ninety minutes, we've covered the seven key issues in security, three techniques for auditing, and two behaviors which routinely compromise security."

- Audience appreciation: "I'd like to thank you all for taking the time to listen. Technical documentation is a hard job and I hope that this training session has made it at least a little easier."

- Next steps: "If you want to know more, I have a page on my site filled with links to excellent articles on the topic. Just visit..." or "There's a professional association and a user group that may be of value to you. They can be reached at..."

- Added value: "Since this is such a complicated topic, I will be available tomorrow for one-on-one sessions to follow up. Just make an appointment with..."

In all cases, be prepared to take questions. I've had teachers who were only a page or two ahead of the class. You never want to be perceived that way. When you are in a situation where people will see you after training — for example, you are an employee or recurring consultant — you should expect that the follow up questions might come days, weeks, or months later. I've had questions as much as five years after an event. Rule 5: *Have a beginning, a middle, and an end*.

The bottom line is that bad training can kill the essential communication of a good idea. IS

CHARLES BAROUCH is the **CTO** for Key Ally, Inc. He is also a Past President, as well as currently a Vice President of U2UG, and a regular Spectrum Magazine contributor.

CHARLES BAROUCH Results@KeyAlly.com
- Consulting; (718) 762-3884x1

Diary of a PDA Project

BY NATHAN RECTOR

I n the March/April issue, Charles Barouch stressed the importance of planning in IT projects. When I read this article, it brought to mind an actual project I was consulting on a few years ago. This project was a PDA project, and illustrates the importance of planning our projects before we start code.

Who? What? When? Where? Why?

One of my long-standing clients called me in regards to PDA project for their production and installation people. This business works in the construction industry making, delivering, and installing products ranging from doors to windows to cabinets to the pretty molding you find around the baseboards of your home. In addition to these areas, they also have a retail hardware store.

This client wanted to do several different things with PDAs. One was to provide a device that their manufacturing production people can enter job status and completion dates as each job is done. Another was to provide their installers electronic time sheets. Another was to provide their delivery people with an electronic copy of their delivery paperwork and locations of the materials to be delivered, then to capture the customer's signatures.

Even though I've mentioned three different projects here, for the initial planning I treated it as one large project with three sections to it. The reason I chose this was because we needed to decide which type of PDA they were going to use. There was no reason start one

project with a Palm platform and find out that another project required a Windows CE platform.

So, to start with, we need to ask the following questions: Who? What? When? Where? Why?

- Who are we creating these programs for?

- What are we going to be doing with the data?

- When is the device going to be used and when it not?

- Where will it be used?

- And last, but the most important, Why? Why do you need the PDA in the first place?

Data Here, Data There!

Let's me start with the first project that deals with the manufacturing floor and answer some questions:

What are we going to be doing with the data? In this case, the user will be updating job statuses and marking which lines on the order are complete. The items manufacturer will be bar coded so during completion bar code will be attached and scanned.

> **" PDAs are nice in production environments because they fit into shirt pockets or can be clipped to belts. There is only one problem with this — they also fall out of pockets or get knocked off belts. "**

Where is the data coming from? The data will be coming from the MultiValue host system.

Where will the data going to? The data will be returning to the MultiValue Host system.

Does the data need to be real-time or can it be batch? This data needs to be as real-time as possible, since customers and salespeople will be able to check this information online. They also need to get updates on the job and questions answered as they come up. Any urgent orders need to be placed in front of them as the day goes by.

Installation Department:
What are we going to be doing with the data? The data is used in payroll.

Where is the data coming from? The data will be coming from the MultiValue host system.

Where will the data going to? The data will be returning to the MultiValue Host system.

Does the data need to be real-time or can it be batch? Since this information will be collected through out the day and is only payroll information, it can be batched and downloaded at the end of the day.

Delivery Department:
What are we going to be doing with the data? They will be using it to view what they need to load on the truck. They need to mark the items as deliv-

ered, to be updated later in the company's system. And they need to capture the signature of the person signing for receive of the material.

Where is the data coming from? The list of items to be delivered will be from the MultiValue host system.

Where will the data going to? The items actually delivered will be returning to the MultiValue host system, but the signature must be stored on a PC file server since it will be a graphics file.

Does the data need to be real-time or can it be batch? This information can be batch-mode since once the truck is loaded, they can't do anything until they empty it and return to get the next order for delivery. The information then can be downloaded and the new order can be uploaded.

You'll notice I have not yet decided on a Window CE or a Palm OS device yet. We still do not have enough information to decide this, other than personal preference which some people will already have going into a project.

Oops. I Dropped It In The Toliet?

The "Who" and "Where" question are very important to the type of PDA hardware device, not to mention the cost of the device, that is chosen.

A lot of the reason PDAs are so nice in production environments is because they fit into shirt pockets or can be

clipped to belts. There is only one problem with this — they also fall out of pockets or get knocked off belts when working in some environments. So the question of "Where" is a very important question and a lot of the time goes hand in hand with "Who".

So let's answer these questions for these projects.

Manufacturing Department:
Who? The person that is going to be using this is a general laborer. They are going to be moving around, lifting heavy things, and using saws and hammers. They are not use to working with computers and are likely to drop the PDA off a forklift, or smash it with a 2x4.

Where? This will be used in a warehouse with a lot of dust and activity.

Installation Department:
Who? The person who will be using this will be in and out of other companies or people's homes installing things. They will be on ladders or under counters. They are likely to leave the PDA at a job site.

Where? It will be used on the road and at job sites where it can be left by accident and/or viewed by people that should not have access to the information.

Delivery Department:
Who? The people who will be using this will be drivers. They are not use

Tech Tips

UniVerse/UniData TOXML

Creating XML documents from a LIST or SORT

statement is really easy in UniVerse 10 and/or UniData 6 with the TOXML sentence option.

TOXML will cause the output from a LIST or SORT sentence to be sent into an XML builder, which reformats your output as an XML document. See figure 1 for an example.

There are a few options that you can use depending on how you want the XML document created. These options are:

ELEMENTS – directs the XML parser to construct the document as "element-centric" (the default is attribute-centric). See figure 2 for an example of Elements output.

WITHDTD – precedes the XML document with a DTD

WITHSCHEMA – precedes the XML document with an XML schema

By default, the XML is displayed as output that can be captured using the CAPTURING clause of an EXECUTE statement. But you can also route the created XML document to a file. To do this, you need to include the "TO *document_name*".

This will redirect the output to a file (the _XML_ in UniData or &XML& in UniVerse) with a record ID of *document_name* with the extension of ".xml" (fig. 3).

If you have specified either WITHDTD or WITHSCHEMA, you will get a second record in the file containing this information. WITHDTD will generate a record with the ID of *document_name*.dtd, whereas the record ID from WITHSCHEMA will use the extension ".xsd".

There is a lot more you can do with the built-in U2 XML parser that we will cover in other articles or Tech Tips. **IS**

Do you have a Tech Tip to share? E-mail it to editor@intl-spectrum.com.

Fig. 1

```
LIST CUST TOXML TO CUST
CT &XML& CUST.XML
```

Fig. 3

Fig. 2 Example of Elements output

DIARY OF A PDA PROJECT

to working with computers, and for the most part are use to getting in and out of the job site as quick as possible. The customer will also be signing their name upon receipt of the material.

Where? The device will be in a truck, or in the rain or snow at the job site.

Now based on this information and the "What" information, we still can isolate some information. We need a device that is rugged, that can be dropped and still work. It also needs to be something reasonably priced since it *will* be broken.

The people that will be using the device are not computer-oriented, so the device needs to be easy to operate and as "idiot-proof" as you can get. Since the devices need to read bar codes, this made us look in the direction of Symbol Technologies (www.symbol.com). They have a device that includes a bar code scanner and is designed to be dropped from a distance of three feet with no damage. Being dropped in the toilet or run over with a forklift is not covered unfortunately, but they are not cheap devices. They start around $1,000.00 list.

This Is Too Disrupting!

The next stop is to work with the people to find out what best works for them when using the device. This is the "When" question. This question usually doesn't come up or cannot be answered during the design stages since it requires the users to work with the software and devices.

The "When" question helps isolate how complex or simple the application becomes. Again, here are the examples:

Manufacturing Department:
When? As each line of a job is complete, the status needs to be entered. This can be 50 items an hour, but on average it is 10 items. Since filling in this information can be disruptive to the produc-

tion process, , the user should have to do very little.

Installation Department:
When? As the person starts each job and finishes each job, they need to update their time and what they where working on: painting, installing, remolding, driving, lunch, and/or break. This is payroll information and is used to pay the installer the correct payroll rate depending on what they are doing. This must be as accurate is possible, and is done at the end of each job as part of the check off list.

Delivery Department:
When? The signature is captured after everything is unloaded, and before the driver leaves.

This Isn't Going To Work.

The following "Why" question should always be asked: "Why do we really need this?" PDAs are the latest and greatest technology to include in a business — make sure there is an actual application for the devices.

During the "Why" questioning, we found the Delivery Departments original idea for using the PDA was not a valid reason. Since hard copies had to be given to the customer of what was delivered and received, they still needed the printed paperwork. Since printed paperwork was still required, having a PDA for loading was not required.

Conclusion

After all this, the choice for the PDA to use was the Palm, due to the size and simplicity. It was the easiest to use for non-computer-oriented people. If the Delivery project had continued, the Windows CE my have been chosen due to a tablet size device needed for delivery to function. The screen size on the palm devices was too small to display all the information needed and capture a signature. **IS**

SHORT-LINKS

What are those short links in the text of some articles?

As our authors share their knowledge and experience, they also point you to the Web for more info. Some of these URLs are long and cumbersome to use.

To make the URL easier to read and use, we are now using Short-Links to redirect you to the web page the author recommends.

The format (intl-spectrum.com/s9999) drops the "http://www." from the front.

Most modern browsers fill that in for you, or you can add it back in.

Comments or additional suggestions? Please drop us a note at editor@intl-spectrum.com

Clif Notes *Continued from page 47*

have to. It does, however, cause grief to IT personnel who try to buffalo their way through something by making up "answers" on the fly or worse, do the condescending "Who are you to question me?" routine.

So one day this dude from IT stops by, and several of them get to talking about the organization's decision to attempt letting them telecommute one day a week. IT dude is going on about how you'll be able to use your home PC to connect to the PC on your desk. One of my daughter's coworker's asks if she will be able to do this from her Macintosh.

"Hopefully, not," IT dude says, turning up his nose. "Besides, why would you want a Mac anyway?"

Now, I guess it's my fault. My daughter seems to have a genetic inability to suffer fools lightly. Noting her coworker's embarrassed silence, she turns to IT dude and says, "Um, because she doesn't want her home computer crashing every time she turns around?"

IT dude laughs and says, "Windows doesn't crash, if you know what you're doing."

If you haven't already done so, please be sure to read Shannon's article "The Top Five Customer Service Mistakes IT Professionals Make" on page 8 of this issue. Look at what this guy did: 1) He refused to say "I don't know, but I'll find out."; 2) He substitutes attitude for competence; 3) He belittles his "customer"; and 4) He directly insults another "customer" in front of her coworkers.

I don't know how many of you have ever been targeted by The Look. If you have, you know that it makes the techniques Jedi Knights employ by harnessing and using The Force look like a bunch of parlor tricks. We still don't know if IT dude is hiding somewhere in the cabling under the floor of the server room or if he was teleported to a synchronous orbit on the dark side of

It's All About the Xusers

BY CLIFTON OLIVER
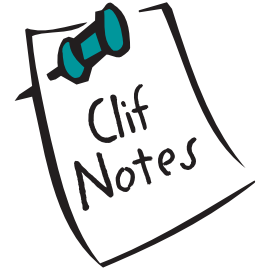
One of the things we do at layout time when producing International Spectrum magazine is select the **"lead"** article and then try to find a cover image that supports that article**'**s theme or topic. We purchase photos or graphics from a number of stock photo web sites. Most of these sites let you enter keywords relating to the subject and present you with a gallery of page after page of images that might fit your need. Sometimes one person does the search and the rest of us voice our preferences. Other times**,** several of us do our own searches and we pool the results. This was one of those times.

Not having seen anything in the preliminary round that particularly grabbed my interest, I decided to do some searches myself. I fired up Safari and pointed it at www.CheapPhotos-sRUs.com (not the real site for reasons that will become evident momentarily), and entered the search keyword "training." Eh. Lots of photos of children in classrooms. So I tried "business training." Well, better. But the photos were mostly of presenters pointing at stock charts, pie charts, and other graphs. Not exactly the kind of training we as IT people do a lot of. So I tried "computer training." That, at least, yielded some photos of people sitting in front of workstations, training classrooms filled with PCs, and a rack of free-weights.

Huh?? (I get very articulate when surprised.)

I am sitting there staring at this photo of a rack of weights used for weight lifting and trying to figure out just what in the world free-weights have to do with computer training. Oh, I could see a link between "weight training" and "training," but not "computer" and "training." And being in the middle of the first couple of pages, it wasn't like this was some images-vaguely-related-to end of listing collection. It was like someone had specifically tagged this photo in order to link it to computers, training, computer users, or... Then it hit me (the idea, not the rack of weights).

Dumbbells. This is another term for free-weights used in gyms. It is also — at least in American English — a slang term for people who are somewhat lacking in intelligence, though not a synonym for politicians. This had all the appearance of somebody loading a picture of dumbbells and deliberately tagging it so it would appear as some kind of statement about people who required training in how to use a computer. Now, I can't state that as a fact, because I don't know. But it has the smell of other stunts I've seen from sophomoric IT dudes with attitudes.

As an example, my daughter works in the accounting department of a large organization that has fairly good sized IT department — an IT department with an attitude, from the stories I've heard. Somewhat unfortunately for a few of them, my daughter has been using computers since she was old enough to sit in a chair in front of a keyboard without having to be propped up by pillows. I didn't teach her. I showed her the mouse, the menus, and the keyboard and turned her lose on her own machine to experiment and learn by doing. I told her there was nothing she could break that Daddy couldn't fix, and got out of the way. (As a side note for you other parents, if you don't mind re-installing Windows, Mac OS, or Linux a couple of times a week, it's a great way to get your kids into computers. At some point, you hand them the CDs or DVD and tell them, "from now on, you break it, you fix it.") As a result, she's a pretty good Help Desk technician, providing Tech Support for her friends and family. This is a Good Thing, if for no other reason than it means I don't