SYS-CON

INTEGRATION • WEB DEVELOPMENT • MIGRATION

WebSphere® JOURNAL

www.SYS-CON.com/WebSphere

The World's Leading Independent WebSphere Resource

AUGUST 2004 VOLUME 3 ISSUE 8

DISPLAY UNTIL OCTOBER 31, 2004

$9.99US $9.99CAN

08

0 74470 03422 3

Will the XP Update Break Your Applications?

BY JACK MARTIN

Microsoft has released an XP update with advanced security technologies and has urged Windows XP users to turn on Automatic Updates to get consistent security updates. However, IBM has told its users not to install the Windows XP update.

If IBM, with close to 400,000 desktops, is not utilizing the XP update, should you? IT departments and developers around the world now face the question of whether or not they should update their systems.

Anyone using or producing Web-based applications (which is all that WebSphere is in its most simple definition) will need to test and, in many cases, modify their applications to work correctly with SP2. Currently, some high-profile, business-critical applications are known to conflict with SP2.

IBM is planning on deploying a customized version of SP2 when and if all of the current issues and concerns have been sorted out and addressed. The question you must ask yourself is, "Can I deploy a customized version of SP2?" Maybe it's time to move your users to Linux or Mac OS X.

Microsoft must be kidding by releasing a security update that disables a significant amount of applications for enterprise customers. In reality, this update is geared at helping the guy at home who uses AOL or another similar ISP and who downloads anything at which his browser is pointed. Maybe there should be a version of XP sold only to Enterprise clients to avoid mixing the two together.

If this was a problem in the automotive world, it would be like Exxon announcing a new type of fuel additive for automobiles. All existing tractor trailers would have to pull over to the side of the road and create a custom mixture for their fuel needs. Otherwise, they'd run the risk that their trucks would stop running.

The Windows XP update will be available soon through downloads, retail distribution, and free CDs, as well as on new PCs. However, nobody knows if it will work, solve the security problems, or render applications incorrectly. A network installation package will be available for Enterprise users who are true masochists.

"Service Pack 2 is a significant step in delivering on our goal to help customers make their PCs better isolated and more resilient in the face of increasingly sophisticated attacks," said Bill Gates, chairman and chief software architect at Microsoft. "It is the result of sustained investments in innovation and extensive industry collaboration. It also reflects a broad recognition that as the security environment changes, the industry needs to work together to respond."

During the worldwide rollout of Service Pack 2, Microsoft will localize the software in 25 languages over the next two months and will begin its distribution. The easiest way for current Windows XP users to ensure that they receive Service Pack 2 when it's released in their language is to turn on the Automatic Updates feature in Windows XP. Customers with Automatic Updates enabled will receive the latest security updates for Windows XP, along with updated installation software that will optimize the download experience of Service Pack 2 and all future updates to Windows XP. The timing for customers to receive the Service Pack 2 download through Automatic Updates depends on a number of factors, including the customer's Internet usage, location, language, and the level of Internet demand for Service Pack 2. Automatic Updates uses spare Internet capacity to progressively download updates without interfering with daily PC use. Microsoft expects to distribute Service Pack 2 to approximately 100 million PCs through Automatic Updates over the next two months.

Jack Martin, editor-in-chief of ***WebSphere Journal***, is cofounder and CEO of Simplex Knowledge Company, an Internet software boutique specializing in WebSphere development. Simplex developed the first remote video transmission system designed specifically for childcare centers, which received worldwide media attention; and the world's first diagnostic-quality ultrasound broadcast system. Jack is coauthor of ***Understanding WebSphere***, from Prentice Hall. jack@sys-con.com

WebSphere® Journal

Advisory Board
Richard Arone, David Caddis, Mike Curwen, Devi Gupta, Lloyd Hagemo, Barbara Johnson, Shannon Lynd, James Martin, Doug Stephen, Victor Valle

Editor-in-Chief JACK MARTIN
Executive Editor GAIL SCHULTZ
Contributing Editor PATTI MARTIN
Product Review Editor JAY JOHNSON
Special Projects Editor RICHARD ARONE
Editor NANCY VALENTINE
Associate Editors TORREY GAVER, JAMIE MATUSOW, JENNIFER VAN WINCKEL

Writers in This Issue
David Lawrence Epstein, Karla Feuer, Glenn Harris, Jeffrey Liu, Cameron Majidi, Jack Martin, Thina Natarajan, Joseph Schwartz, Raman Sud, Helen Wong

Subscriptions
For subscriptions and requests for bulk orders, please send your letters to Subscription Department.
SUBSCRIBE@SYS-CON.COM
Cover Price: $8.99/Issue Domestic: $149/YR (12 Issues)
Canada/Mexico: $169/YR Overseas: $179/YR
(U.S. Banks or Money Orders)
Back issues: $15 U.S. $20 All others

PRESIDENT AND CEO FUAT KIRCAALI
PRESIDENT, SYS-CON EVENTS GRISHA DAVIDA
GROUP PUBLISHER JEREMY GEELAN
SENIOR VP, SALES & MARKETING CARMEN GONZALEZ
VP, INFORMATION SYSTEMS ROBERT DIAMOND
VP, SALES & MARKETING MILES SILVERMAN
PRODUCTION CONSULTANT JIM MORGAN
FINANCIAL ANALYST JOAN LAROSE
ACCOUNTS PAYABLE BETTY WHITE
ADVERTISING DIRECTOR ROBYN FORMA
ADVERTISING MANAGER MEGAN MUSSA
ASSOCIATE SALES MANAGERS KRISTIN KUHNLE, BETH JONES, DOROTHY GIL
LEAD DESIGNER LOUIS F. CUFFARI
ART DIRECTOR ALEX BOTERO
ASSOCIATE ART DIRECTORS RICHARD SILVERBERG, TAMI BEATTY, ANDREA BODEN
WEB DESIGNERS STEPHEN KILMURRAY, MATTHEW POLLOTTA
CIRCULATION SERVICE COORDINATORS EDNA EARLE RUSSELL, LINDA LIPTON

Editorial Offices
SYS-CON Publications, Inc.
135 Chestnut Ridge Road, Montvale, NJ 07645
Telephone: 201 802-3000 Fax: 201 782-9637
SUBSCRIBE@SYS-CON.COM
WebSphere® Journal (ISSN# 1535-6914) is published monthly (12 times a year).
Postmaster send address changes to:
WebSphere Journal, SYS-CON Publications, Inc.
135 Chestnut Ridge Road, Montvale, NJ 07645

WORLDWIDE NEWSSTAND DISTRIBUTION
CURTIS CIRCULATION COMPANY, NEW MILFORD, NJ

FOR LIST RENTAL INFORMATION:
KEVIN COLLOPY: 845 731-2684, KEVIN.COLLOPY@EDITHROMAN.COM
FRANK CIPOLLA: 845 731-3832, FRANK.CIPOLLA@EPOSTDIRECT.COM

Two years without a vacation.
The application's up. It's down.
It's up. It's down.

I'm to blame. Steve's to blame.
Someone's always to blame.

Not any more.

Get Wily.™

Enterprise Java
Application Management
1 888 GET WILY
www.wilytech.com

wily technology

TOM INMAN
VICE PRESIDENT OF PRODUCT MANAGEMENT AND MARKETING, IBM WEBSPHERE SOFTWARE

What's Happening in WebSphere

A look at what's in store

Last month, Jack Martin, editor-in-chief of *WebSphere Journal*, and Tom Inman, vice president of product management and marketing, IBM WebSphere Software, talked about the differences between WebSphere and WebLogic. This month, they look at what's happening in WebSphere now, and plans for the future.

Jack Martin: What was your background? What did you do before you did this?
Tom Inman: I joined sales and marketing right out of college and have been here ever since.

JM: What school did you go to?
TI: I went to Michigan State University, where I earned a Bachelor of Science in engineering and an MBA in marketing and finance. I needed money to pay for school and I landed an internship with IBM just outside of Michigan State in Lansing, Michigan. I hit it off with the branch manager and the sales manager very well. They tried to hire me before I finished school and they were going to pay for the rest of it. I thought the opportunity to work for IBM made sense, but I wanted to get out of Michigan, having grown up there. It is a great place to call home, but I was young and needed to explore more of the world. I coerced them into helping me find a job in San Francisco because I wanted to move to the West Coast. This was back in the mid-'80s.

I spent a number of years in San Francisco, but I have been here in the New York area for 8 years.

I have spent my career mostly between sales and marketing. I had a short stint as the executive assistant to Dr. Irving Wladawsky-Berger, IBM's head of technology and chief visionary.

JM: I've met him twice. They were both enlightening conversations.
TI: He is a visionary who actually is quite grounded in reality. In my experience working for Irving, there was some technology developed by IBM for the 1996 Olympics and the Web site. Irving asked me to take a look at this technology to see if we could actually make a business out of it. To make a long story short, the ideas and technology used for the Olympics became the seeds for what became WebSphere. We set up a small group of people with business and technology backgrounds and built the initial business plan.

JM: Were you involved when Don Ferguson first got involved?
TI: Yes. Don was working on an initiative, which I'm sure he talked to you about, called IBM Component Broker. Don's ideas, which went into Component Broker, in many ways became the reference architecture and design that shaped most of the major technology specifications that went into the creation and development of J2EE.

There were some great ideas behind Component Broker and some challenges. One of the challenges, quite frankly, was that it was too complex for what the general market was seeking. I pushed very hard for a more basic, simpler technology and product to introduce into the marketplace.

At that point, the predominant customer need was for a simple, standards-based Web interaction server. This would enable self-service applications to be built and deployed on the Web and would enable connectivity to some back-end systems and applications such as CICS, DB2 databases, Tuxedo, Oracle databases, or whatever back end the customer might have. In that, we needed a product that was properly suited for the market and so I helped push the design for what became the first WebSphere products. Those were WebSphere Application Server and WebSphere Studio, which we introduced in May of '98.

So, did I work with Don? Absolutely. And we have worked together to translate various engineering ideas into appropriate offerings designed to meet the needs of the market over the years.

JM: I didn't know that you were one of the original WebSpherians until just now.
TI: I like to keep a pretty low profile.

JM: I've heard from a few people that you're an enormous advocate for your clients. How do you go about actually advocating, inside an enormous company like IBM, for a specific customer?

TI: That's a good question, Jack, one that I appreciate because our clients' needs and their satisfaction with IBM as a provider of technologies and solutions is why IBM exists. My first priority is to spend enough time with the clients and the business partners to learn first hand what it is they need. The good news about customers and business partners is that they're pretty vocal when they have a problem they're trying to solve. They are not shy about telling people about it.

I think those companies that can actually step forward and solve their problems are those that are going to get the business. So, lesson number one is to listen closely and, especially, to understand what it is that the client is seeking to do.

I use my position in the organization to be an advocate for the client's needs. I work closely with the product managers, the lead architects, and the lead development managers on those particular requirements to make sure that what we are creating meets the needs of the real world clients.

JM: Are you saying that if you had a customer that said, "Tom, listen, we've got this problem and it's Problem X" and the product didn't currently support it that you could potentially have those types of features brought forward?

TI: Absolutely. We do that regularly. That is, we ensure that we are listening to what the client needs and we ensure that we're confirming our understanding. We also need to be careful that we're addressing the needs of the general market and not simply the needs of a single customer. In the software products business, we're trying to build products that capture 40+% market share, like our WebSphere Application Server business, or 70+% market share like our WebSphere MQ messaging portfolio does. To capture this level of market share, you need to focus on client needs that represent the general market and then do a better than your competitors do in meeting the clients needs.

JM: Repeatability of sales is key.

TI: Repeatability of sales is key, so you need to make sure you are addressing the more general marketplace. Having said that, if you can find the early adopter customers that reflect the general market and partner with them, that's a very good way to accelerate delivery of products and we do that regularly. If you go back to the very first seeds of WebSphere, we put an engineering team in place in a financial services customer in the Northeast so that we made sure that what we were building reflected their needs as opposed to what we thought we might need at the development lab.

The same holds true for the first versions of the portal product. We had a good prototype, we had good designs, we said to our clients, "Okay, talk to us. Let's get out of the lab and go solve a real problem within the market." We rolled up our sleeves with a large automotive company and helped them establish an employee portal and then a customer and supplier portal. Again, we did that because we knew that their problems represented those of the broader marketplace. We figured if we applied the technology to their problem, we'd accelerate our learning. We're doing that again today around the whole space of Web services, helping clients to design and deploy their service-oriented architectures. We are rolling up our sleeves with clients and using the experience to help find and then iteratively fill the gaps we might have in our products. We use these experiments to help develop capabilities and then grind those requirements back into the core engineering and product development process. It helps us get our offerings right and accelerate time to market. I call it market development.

JM: I knew that IBM was doing that. I didn't realize that, at the sales level, someone like you was getting involved. I thought that was more of a research and development area.

TI: You know, much of the understanding comes from a background in sales, but my current role is vice president, product management and marketing, WebSphere. In the product management area, the responsibility is, in fact, to help craft the product portfolio. Working with various development teams across the various products and the portfolios to ensure that we are building the right products and then bringing them to market. As you might expect, clients typically need more than one of those products, so the question becomes how to make the client's experience with our portfolio exceed the value of the sum of the products.

One of the things that I'm very intently working on now is how we ensure that the products in the client scenarios that we're targeting, those where the client needs span more than one part of our portfolio, are well inte-

"our clients' needs and their satisfaction with IBM as a provider of technologies and solutions is why IBM exists"

"We will enable mid-market companies to realize profit from enabling business flexibility and responsiveness"

grated and work well together to solve actual problems. So, it takes some sales background, with understanding clients' wants and needs, and then translating that into product road maps and product development and product packaging decisions and that's what I spend a majority of my time on.

JM: You're still regularly out there seeing customers?
TI: Absolutely. Last week, we held our WebSphere Inner Circle event, which is our top 200–300 customers and we had them together for three days in Las Vegas. At the Inner Circle we host a three-way information sharing – IBM to our clients, our clients to us, and our clients with each other. Topics such as best practices, what works, what doesn't work, feedback, you know, about our product road maps. I recently spent time with many of our top customers and, at the same time, we also had about 40 of our top ISV partners out there. We held different meetings, but with a similar structure, sharing our direction with them and capturing their direction back to us, and providing a forum for them to share their experiences with each other.

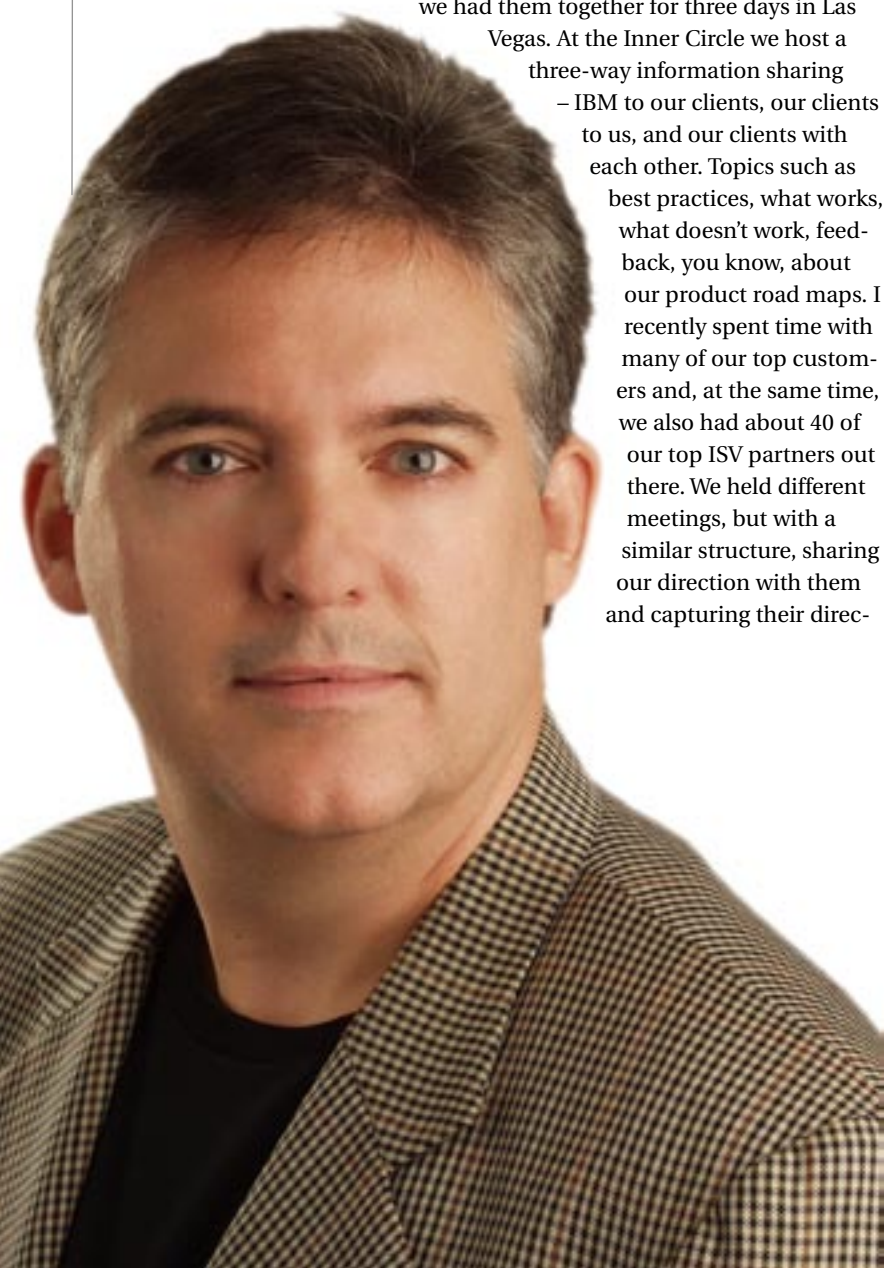

JM: Was it right around then that you announced a new spec that IBM just achieved?
TI: Yes. IBM announced and published its results for something called the SpecJ AppServer 2004 benchmark. We announced our support for that as well as being the first vendor to deliver published results that, by the way, are very impressive.

JM: Could you explain exactly what this performance benchmarking is and what it means?
TI: First of all, in order to help the marketplace make informed technology decisions, the industry often agrees to measurements, or vehicles, to be able to compare and contrast offerings. One of the important characteristics or attributes that clients expect out of their technology infrastructure, to support their business, is performance. The industry mobilizes around performance benchmarks. SpecJ AppServer 2004 is a benchmarking vehicle to be able to compare performance models and metrics across various application server infrastructure technologies. What we like about it at IBM is that we think it properly reflects the typical customer application patterns rather than some artificial situation. It reflects this better than its predecessor environment and forms a decent performance model for clients to use when evaluating performance as a variable in the technology decision.

The performance benchmark was collaborated on by multiple parties across the industry to reflect a real world application. The fact that IBM was the first to announce performance benchmarks, and with very impressive results, is a testament to IBM's leadership in performance.

JM: Where can we expect WebSphere to go in the next 12 months – from your perspective?
TI: The first thing to recognize is that what guides us is our clients, and our clients and our business partners' needs. More and more what we are finding are clients needing increased business flexibility and responsiveness. That requires increased IT flexibility and responsiveness. So, first, you will see us continue to build out the principles of a service-oriented architecture continuing on from the announcements we made in April. This will include enhanced technology and enhanced best practices and services offerings.

We'll also be showcasing more and more proof points. Clients value seeing others adopt our technology so that they can learn from their experience. Second, more and more of the business community, as opposed to the IT community, is taking a greater role in overseeing the direction in where their business spends their money on IT.

JM: That's absolutely true.

TI: Therefore, you'll see us do more and more to position our capabilities to line up with the business problems that we see, industry by industry. Such as, how does WebSphere help enable and streamline a supply chain? If that supply chain is in retail, we help track item movement from the point of sale all the way back through the supplier network. This is called item synchronization and is an area where we, along with our partners, already have a solution. Another area is how we help telecommunications clients streamline their business support systems and their operational support systems. We help them implement an OSS/BSS integration hub to enable the integration of the various applications and business processes required to streamline their operations and help them to be serve the needs of their clients.

So, you'll see us deliver more and more solutions and solution building blocks to solve the problems that the line of business has. This gets us back to the prior question of the business partners and their role. You will see our application ISV partners playing a role here in terms of their applications being "integration ready" or their having pre-built adapters or portlets into their environment. Doing so enables their products to be more easily and cost effectively integrated with our applications, systems and processes in an enterprise or a medium business

The third area in which you will see WebSphere investments is around the tremendous momentum we have built in the mid-market by continuing to build out our Express offerings and to build out the supporting partner ecosystem. With this investment, we will enable mid-market companies to realize profit from enabling business flexibility and responsiveness from their use of our WebSphere software. Recently we introduced WBI Server Express to enable process and internal application integration for mid-market companies. This offering was really built from the ground up to meet the needs of mid-market customers and the needs of their solution providers.

The bottom line is that we will invest in continuing to lay out the building blocks and best practices for SOA enablement, to generate more and more solution capabilities, along with our business partners, that are relevant to industry-specific business problems across all major industries, and to continue to build out our relevance in the mid-market by way of Express offerings and in the thriving ecosystem of partners we have there.

JM: So, we've got a lot to look forward to. The best is yet to come!

TI: Absolutely.

JM: What do you see with the emergence of a new market in open source with JBoss?

TI: You need to be careful in equating JBoss with any of the other popular open source efforts out there. JBoss is a company that is taking standards-based application server software and introducing a different business model around it. The business model is to make the software "free" to the "purchasers" of the software and then to offer service and support around that free software. So, it's really a different business model for a commercial software company that's actually going to make its money off of service and support licenses. Services for a fee.

Do I think there's a place for a model like that in the marketplace? Yes, I do. I think it will appeal to folks who would like a lower-cost point of entry from such technology, but will require a level of service and support poured around that. It's questionable, in my mind, how long JBoss can remain that model provider because they are a very small business. They don't have a tremendous reputation for being world-class in the area of scalability or service and support, unlike somebody like IBM with a strong brand identity for service and support.

Do I see them gaining some share? Yes, I do. I see them gaining share, frankly, largely at the expense of somebody like BEA.

JM: Why is that?

TI: Because, as you know, we at IBM are just gaining the overall momentum. Second, when a buyer thinks service and support in the IT industry, one of the first companies they think about is IBM. BEA doesn't necessarily have a long-standing reputation for being capable of providing world-class service and support, and neither does JBoss. I should also note that JBoss will be seeing some competition in the open source area once Apache Geronimo comes along a bit more, so we'll see how it weathers that.

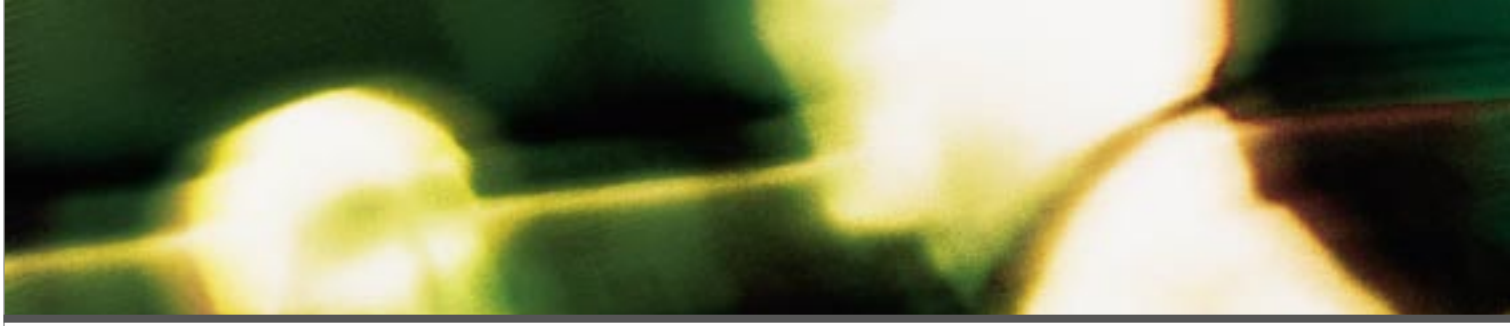

Business Intelligence

Jack Martin also spoke to Dave Lai, CTO of Viador

JM: What does Viador do?
DL: Viador is a business intelligence tools ISV.

JM: What type of Business Intelligence tools?
DL: We provide reporting and OLAP analysis tools for customers that need to do reporting and analyze their data such as financial analysis.

JM: What customer base do you have; who uses your tools?
DL: Our customer base is quite broad. We have Fortune 500 companies like Citibank, Verizon, JP Morgan; and we also have medium-sized shops like Spectra Marketing, which provides a hosted solution (ASP). Viador also has a very strong presence in China. We are one of the largest business intelligence tools companies in China right now in terms of the support infrastructure.

JM: So you do business all over the world?
DL: Yes. Actually, we have been in business since 1996. We have customers in China, Japan, Europe, Russia, United States.

JM: Your product works with an application server, correct?
DL: Yes; it can be integrated with WebSphere Application Server natively.

JM: Does it work with any other application servers?
DL: Right now, our strategy is focused more on the WebSphere application server, from both the technology and the business perspective.

JM: Why is that?
DL: From the business perspective, we have a long history with IBM – since 1996, when the Internet just started. We worked with IBM to deploy a lot of solutions and we have a joint customer with IGS. Recently, in 2002–2003, we elected to integrate Viador technology with IBM WebSphere technology natively. Actually, a lot of our customers want to extend the usage of WebSphere technology, they want something that integrates into their existing environment instead of a proprietary application server like most of our competitors. That occurs when a customer already has WebSphere and they are looking for a better return on their investment by leveraging the application server they have.

JM: Why do you think most of them have WebSphere as opposed to, say, BEA products?
DL: From the perspective of our customer base they are looking for solutions that go beyond just an application server. They are looking for a vendor that can offer them a wider spectrum of technologies. They brought DB2 and WebSphere together, so it becomes a very strong driving position to our customers. Some of our customers use WebSphere instead of BEA WebLogic because the product offering from IBM is much broader, it goes much beyond an application server. Some of our customers actually have concerns as to whether BEA will be a viable stand-alone company.

JM: So you have had customer concerns as to whether BEA is going to be around either as an independent company or at all?
DL: Exactly. So if they need more than just an application server to run their organization, they prefer a vendor that offers more than just an application server. IBM offers more of those things, like the database management, messaging systems.

JM: Have you had any customers look at WebSphere and WebLogic head-to-head and were you involved in it?
DL: All of my encounters have been to look at that in the U.S. market.

JM: How about over in China? What is the perception of WebSphere and BEA in China?
DL: BEA is perceived as the market leader in application servers over there, but there is a growing trend that customers are looking for a complete integration package rather than just an application server, which IBM offers.

JM: That's the fastest growing information technology country on Earth.

DL: Yes.

JM: What is the impression in China and Hong Kong of WebSphere and WebLogic? Let's look at them one at a time. When people think about WebSphere in Asia, what do you hear them saying?

DL: In China BEA entered the application server market much earlier than IBM. However, WebSphere comes in very strong as well; IBM is investing a lot of resources to push the WebSphere offering, including such things as training, seminars, etc. Customers are getting more and more interested in WebSphere, especially IBM's focus on solution offering really align with customers' interests.

JM: Since you have a unique view of what's going on in China, where do you see the application server market going in the next one to three years?

DL: IT in China pretty much is endorsing the application server. They like the fact that they can leverage the application server scalability, load balancing, security, etc. Now, they can focus on their own application and domain expertise. Moving forward, customers are looking for a total solution, so the vendor that can create an ecosystem that provides a complete solution will be the winner in this market. The key success factor is how each vendor can get quality partners to fill their solution gap.

Another trend in China is Linux. The fact that IBM is pushing Linux big time is a favorable factor for them. Some Chinese companies even have a preference for Linux as an operating system.

JM: Why do you think Chinese people prefer Linux over any other operating system?

DL: This is again a very interesting thing. I think the government and the IT industry is pushing for Linux. With Linux, China now has a different angle to penetrate the IT industry, not just as a consumer, but as a producer. Another perception is that Linux does not have IP issues, people associate Linux with open source. So, that solves the perception issues on IP.

JM: Whose perception?

DL: I mean the buyer's perception with respect to the IP right. In general, Linux is associated with open source and they can develop something on top of it.

JM: So they can do what they want with Linux, where, say, with Microsoft, they can't.

DL: Right.

Streamlining the development process

Surfacing a Simple Domino Application in WebSphere Portal

BY GLENN HARRIS

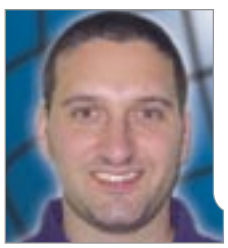

Glenn Harris has been an IBM technology advocate for over a decade. He holds a Master's Degree in Information Systems and has maintained his PCLP since v3.x. Glenn focuses on application development, enterprise integration, and project management. He is currently employed by Solutions By Design II (www.sbd.com). Off-line, Glenn spends much of his time with his wife, Jenny and his two children, Emily and Sarah. Glenn resides in the New York metro area and can be reached at glenn.harris@sbd.com.

Look at it this way, saying surfacing is an improvement over saying portalizing! Just in case this is your first time hearing either term, they have identical definitions: bringing the data to the top most layer of the system. If you think about your development experience, you've probably been surfacing applications for years now. In this article, we'll cover a step-by-step guide to help you surface a simple Domino Content Management application on top of WebSphere Portal using the Bowstreet Portlet Factory for WebSphere 5.8 tool.

You have several options when you consider surfacing a Domino application into your portal environment. In the past, Solutions By Design II has used the Portal API, WebSphere Portal's Portlet Builder for Domino, and Bowstreet's Portlet Factory for WebSphere. Through experience, we have found that development with Bowstreet's Portlet Factory for WebSphere offered the greatest flexibility, streamlined our development process, and didn't involve any coding.

The Domino application provides authorized users with the ability to post a document attachment for the portal community to view. At a minimum, you'll need two fields: a Text field to allow your authors to identify the attached document and a Rich Text field to store your attachments. As seen in Figure 1, our Posting Form uses the new Rich Text Lite feature in Domino 6.5 to ensure that users post a file attachment only in the Rich Text field. Remember, this is a simple database, we would expect you to take your development experience and add extended features like validations, posting URLs, multiple attachments, automatic content archiving/expiration dates, etc. *Note:* Make sure that the User name and Password you will be entering in later steps has proper access to the application, that the Domino Server allows browsing of databases, that DIIOP is running as a server task, and that you have created a full-text index.

You'll also need to create a view for the Portlet to reference. The view should have two columns, one with the document Title and the other with the URL to the attachment. @URLEncode is a new Domino 6.5 feature, which we used in our URL formula.

```
"http://" +
@Name([CN];@ServerName)
+"/"+ @WebDbName + "/0/" +
@Text(@DocumentUniqueID) +
"/$file/" + @URLEncode("Platform";
@Subset(@AttachmentNames; 1));
```

Creating the Portlet with Bowstreet's Portlet Factory for WebSphere

STEP 1 – CREATING A NEW PROJECT

a) Launch WebSphere Studio.
b) Select File – New – Project – Bowstreet – Bowstreet Web App Project, Select Next.
c) ***Project Name:*** Type in a Project Name (I used SBDIIPosting), Select Next.
d) ***Add Features:*** Select Lotus Collaboration Extension, Select Next.
e) ***Java Settings:*** Select Next.
f) ***App Server:*** Fill in the field information, Select Next.
g) ***Portlet Settings:*** Under the WebSphere Portlet WAR section, select Create. WebSphere Portlet WAR, fill in the fields for this section, Select Next.
h) ***Summary and Important Details:*** Select Finish.

Note: An "Output Directory" error will be generated. Save the Project and Deploy the EAR file to your WebSphere Application Server. Once you've deployed the EAR file, return to the project and rebuild it; the error should no longer be present.

STEP 2 – CREATE A NEW MODEL

a) Open the Bowstreet Perspective, Window – Open Perspective – Other – Bowstreet Models, Select OK.
b) Select on the Project you created in Step 1.
c) Right-click on the Project, New – Bowstreet Model.
d) ***Available Projects:*** Select the Project you created in Step 1, Select Next.
e) Select Main_and_Page, Select Next.
f) Enter the Name of your Model, Select Finish.

STEP 3 – ENSURE BOWSTREET MODEL VIEWING

a) Select Window – Open Perspective – Bowstreet Model.

STEP 4 – ADD A PORTLET BUILDER TO THE MODEL

a) In the Outline, there is a Plus icon above the main Action List and the page1 Page. Click on the plus symbol to Add a Builder Call to the Current Model.
b) Select Portlet Integration – Portlet Adaptor, Select OK.

STEP 5 – CONFIGURE THE PORTLET ADAPTOR

a) Fill in the Name and Portlet Title fields, Select OK.
b) Select OK.

STEP 6 – ADD A DOMINO VIEW & FORM BUILDER TO THE MODEL

a) Click on the plus symbol to Add a Builder Call.
b) Select Collaboration – Domino View & Form, Select OK.

STEP 7 – CONFIGURE THE DOMINO VIEW & FORM

a) Fill in the Name, Host server, User name for generation and Password for generation.
b) Click the "Get databases and views" button.
c) Click the "…" button at the end of the Database name field and select the Domino Application you are surfacing, Select OK.
d) For this example, leave Runtime credentials as "Use regen credentials specified above."
e) Select the View name from the drop-down box.
f) Click Paged data display checkbox.
g) Enter 15 or any value you like for Rows per page; this will display paging options when appropriate, based on your inputted number.
h) Click Show search button checkbox.
i) Select OK.
j) Delete the main & page 1 builders from your model.

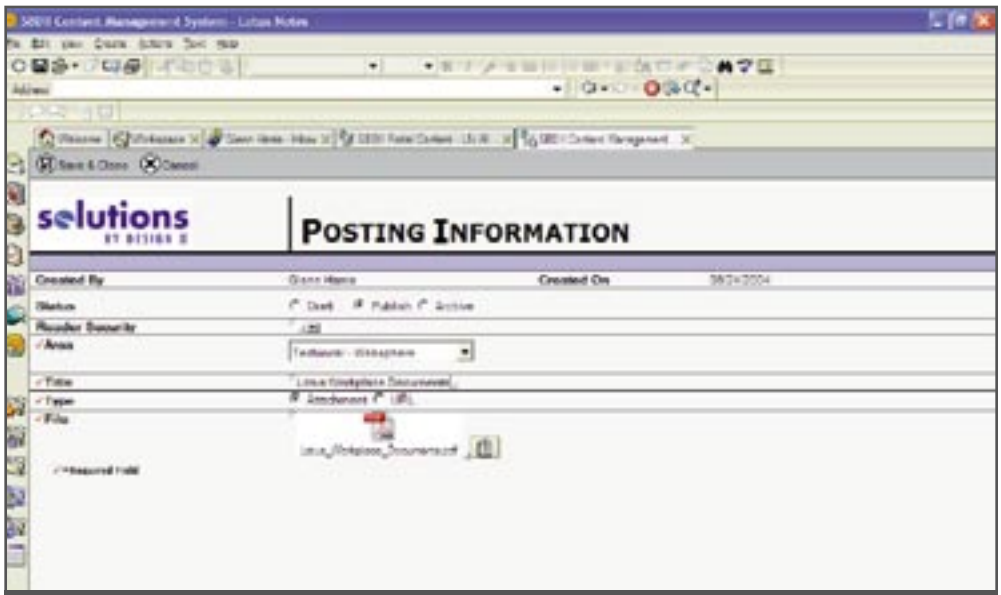

FIG 1: POSTING FORM

FIG 2: CHOOSE PROJECT

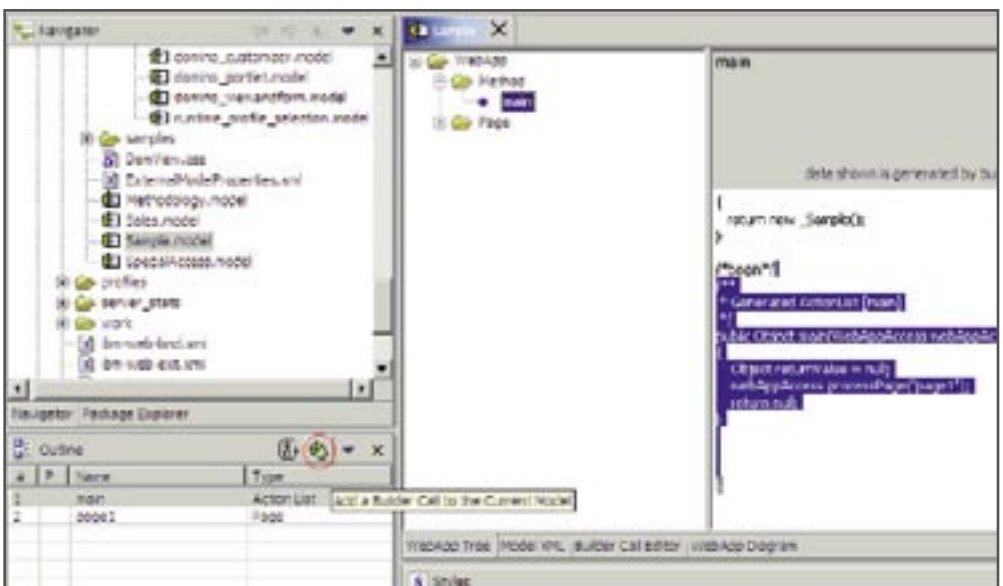
FIG 3: ADD BUILDERS

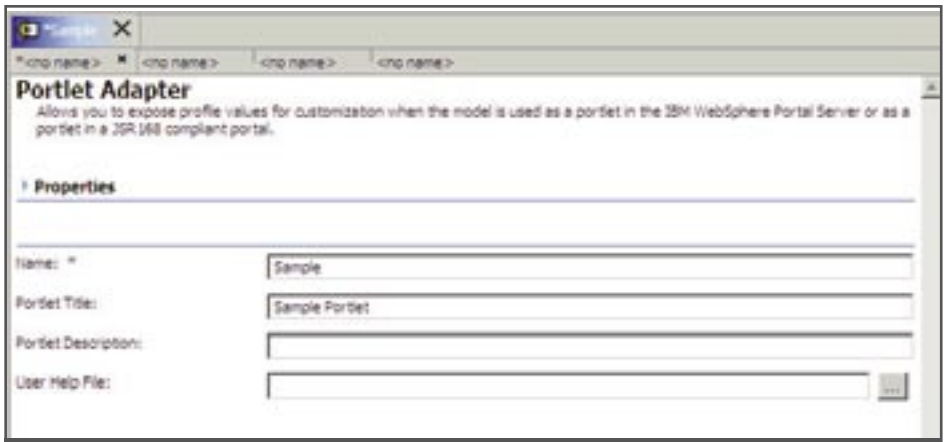

FIG 4: PORTLET ADAPTER

FIG 5: BUILDER DOMINO VIEW & FORM

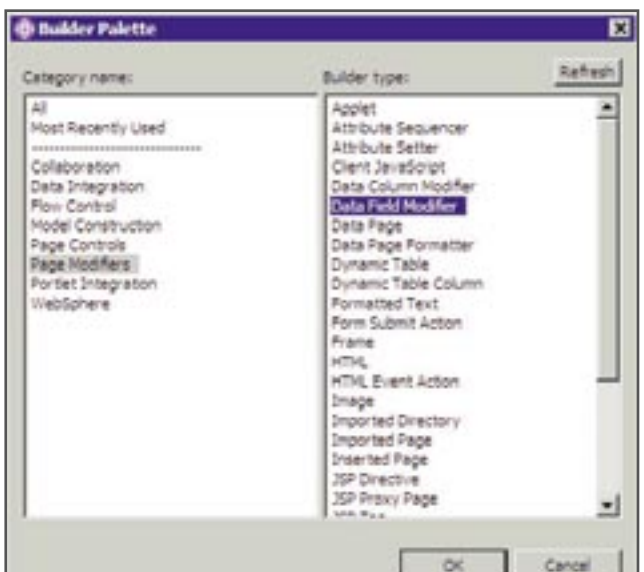

FIG 6: BUILDER DATA FIELD MODIFIER

STEP 8 – ADD A DATA FIELD MODIFIER TO THE MODEL

a) Click on the plus symbol to Add a Builder Call.
b) Select Page Modifiers – Data Field Modifier, Select OK.

STEP 9 – CONFIGURE THE DATA FIELD MODIFIER

a) Fill in the Name.
b) In the Fields section, select the "…" button. Expand the [DataViewPage]DataPage, Expand ViewData, Expand Row and locate your Column name that holds the Title information. In this example, my Column is named Links, Select OK.
c) In the Field Settings Section change the Field Behavior to "Action Field."
d) Change the Action Type drop-down box to "Link to a URL."
e) Select the "…" button at the end of the URL field, Expand Variables, Expand RowLoopVar, Expand Row, and select the column in which your Attachment Link is located. In this example the column is called urladdr. Select OK.
f) Select the Target drop-down box value of "_blank" (this will launch the attachment into a new window).
g) Select OK.

STEP 10 – ADD A DATA COLUMN MODIFIER TO THE MODEL

a) Click on the plus symbol to Add a Builder Call.
b) Select Page Modifiers – Data Column Modifier, Select OK.

"We would expect you to take your development experience and add extended features"

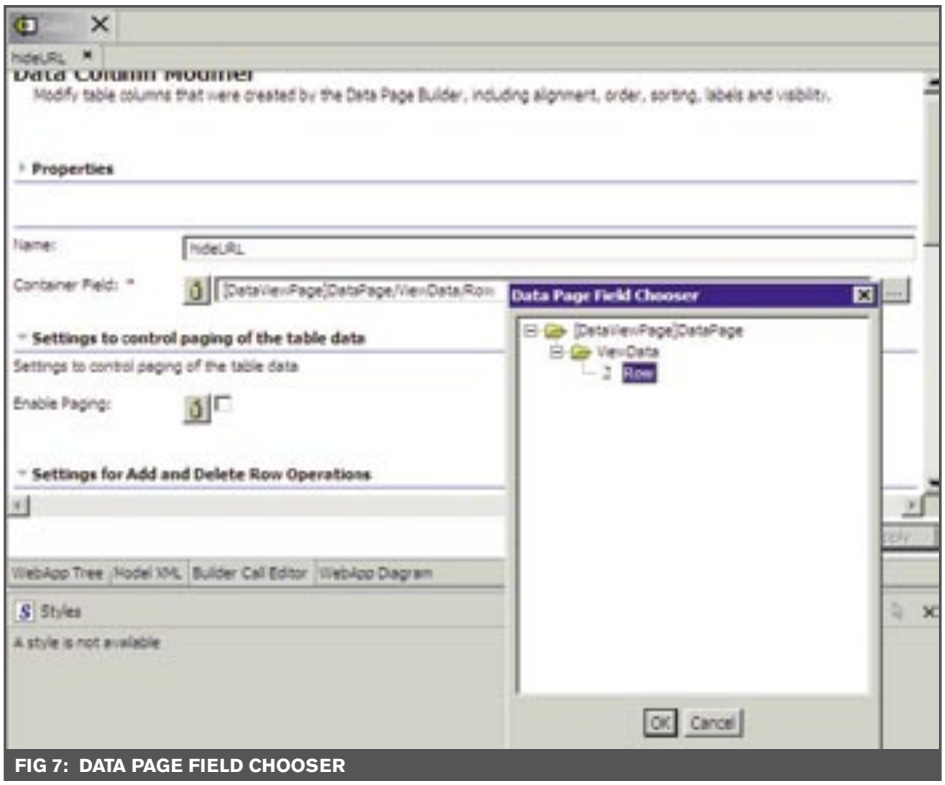

FIG 7: DATA PAGE FIELD CHOOSER

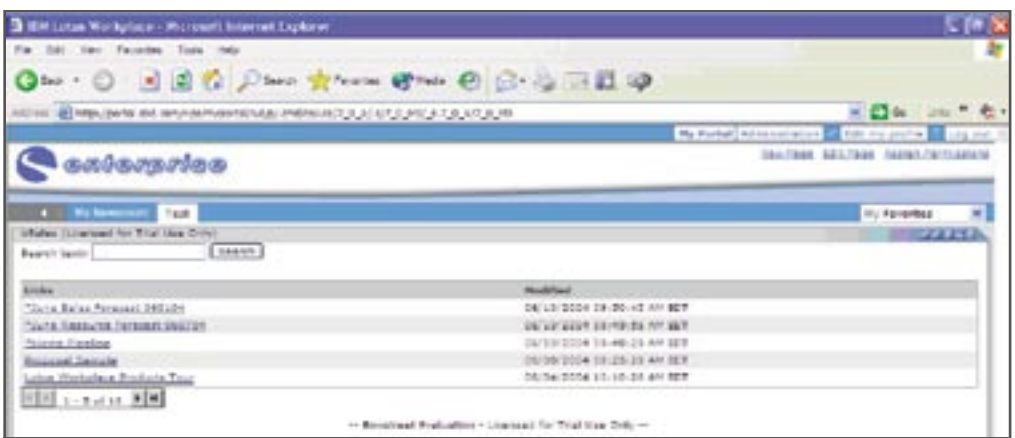

FIG 8: WHAT YOUR PORTAL SHOULD LOOK LIKE

STEP 11 – CONFIGURE THE DATA COLUMN MODIFIER

a) Fill in the Name.
b) Select the "…" button at the end of the Container Field, Expand [DataViewPage]DataPage, Expand ViewData, select Row, Select OK.
c) Scroll to the bottom of the builder. Under Status in the "urladdr" row select Hide.
d) Select OK.

You're now ready to deploy your WAR file to your Portal Server. There's a good deal of documentation for deploying the EAR to your WebSphere Application Server and your WAR to your Portal Server.

Conclusion

You have just surfaced your Domino application inside of WebSphere Portal – congratulations! Your portal should look something like Figure 8.

References

- *IBM developerWorks:* www-136.ibm.com/developerworks/websphere
- *WebSphere Portal Information Center:* http://publib.boulder.ibm.com/pvc/wp/502/ent/en/InfoCenter/index.html
- *Bowstreet:* www.bowstreet.com

"The Domino application provides authorized users with the ability to post a document attachment for the portal community to view"

How WebSphere BI Server creates a hub-and-spoke architecture

At the Heart of an Enterprise Part 2

BY DAVID LAWRENCE EPSTEIN & CAMERON MAJIDI

David Lawrence Epstein is an independent contractor working with Prolifics providing WBI solutions, as well as WebSphere Business Integration Server training courses. He has been designing and developing production software for various companies for 20 years. David has trained managers and developers in Java, EAI, and all areas of WBI Server. He offers self-paced and live Java, WBI Server, and WBI Server Express courses and materials at www.wbiguy.com. david@wbiguy.com

This is the second in a two-part series of articles introducing the benefits of a hub-and-spoke architecture and IBM's WebSphere Business Integration (BI) Server. In the first article, we detailed the financial and technical benefits of using a hub-and-spoke design, such as reduction in software maintenance costs, overall simplicity compared to point-to-point programs, ease of general administration, and a more modular design allowing the business logic to be application independent.

Here, we start with the application independent business process modules contained in the hub, which work their way through the design and purpose of each of the major components of IBM's WebSphere BI Server. After an introduction to the major components, we conclude with the details of handling updated data using WebSphere BI Server cross-referencing tools to synchronize data.

WebSphere BI Server: ICS and Repository

Our focus on the specifics of IBM's EAI offering, WebSphere BI Server, begins with the hub itself. The hub is called the InterChange Server or simply, ICS. It's a Java program and, while running, it stores whatever it needs in memory so it can perform quickly. The ICS uses a database named the Repository for persistent storage; DB2, Oracle, or Microsoft SQL Server can be used to host the Repository (see Figure 1). Should there be any hardware or software failures on the computer hosting ICS, data that was in memory can be recreated, based on the data stored in the Repository.

Recall that in a hub-and-spoke architecture, all integration data flows through the ICS hub, making it a convenient location for integration business logic. As the container of business process modules, the ICS provides an opportunity to separate business logic software from the complexities of application connectivity.

WebSphere BI Server: Collaborations

In WebSphere BI Server, business logic is implemented in components called collaborations. They're application-independent modules, visually shown and developed as flowcharts in a tool called the Process Designer.

Because collaborations are application independent, they can be shared and reused. IBM provides over 70 prebuilt collaborations, potentially deployed with 10 minutes of configuration. These prebuilt collaborations represent hundreds of years of experience working directly with clients – they're the common

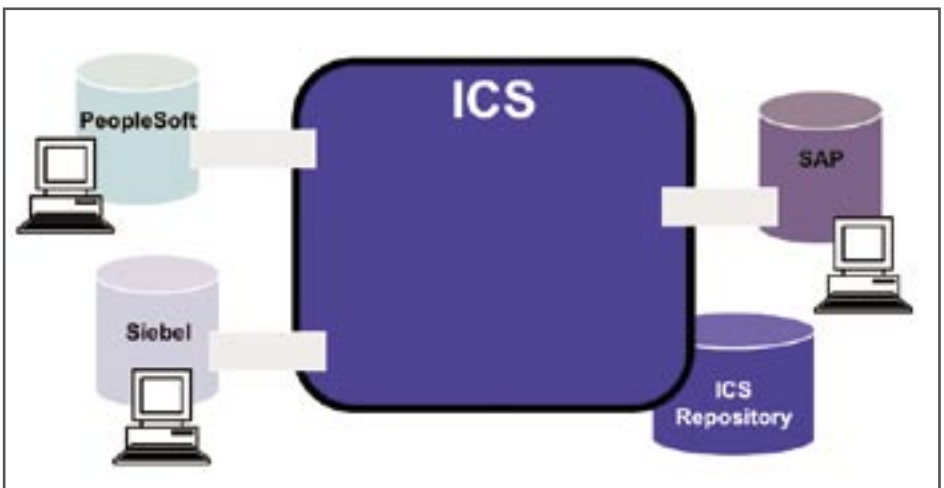

FIG 1: USING ICS AS THE HUB AND ICS REPOSITORY FOR PERSISTENT STORAGE

solutions to real-world problems. Examples of popular prebuilt collaborations include Sales Order Processing, Item Synchronization, and Customer Synchronization.

Given prebuilt business logic modules, as diagrammed in Figure 2, the task becomes one of understanding the logic and knowing what configuration options to set.

WebSphere BI Server: Business Objects

The data on which collaborations operate is packaged in Business Objects (also referred to simply as BOs), which are containers for application data. For example, a business object might hold an SAP customer or a Siebel customer.

Each business object has a name, so it's easy to tell whether it's an SAP customer or a Siebel customer. Each business object also has a verb, such as Create, Retrieve, Update, or Delete. Finally, each business object is full of attributes, such as FirstName and LastName, and their values.

It's important to note the effect that the verb has on the data. Without the verb, the data in a business object represents an entity, nothing more. With the verb, a business object is actually an *event*. Instead of simply Jane Doe, a business object is actually the *creation* of Jane Doe, or an *update* of Jane Doe.

When a business object that represents the creation of Jane Doe is delivered to a collaboration, the collaboration typically needs to know specific information that the business object contains, such as Jane Doe's zip code. In order to obtain specific values, the collaboration needs to know a description of the business object. The data representations of business objects are supplied in Business Object Definitions (also referred to as BODs), stored in the Repository (see Figure 3).

Essentially, a business object is a data structure. At design time, defining a business object consists of naming a series of fields (called attributes) and their data types, and also defining a list of supported verbs. At runtime, each business object instance will have one of the verbs from its list, and will have its attributes filled with actual data. Business object attributes don't have to be primitive data types, they can be other business object types or arrays of other business object types. Because business objects can be nested, they can represent very complex data schemas.

Business Objects and Collaborations

Because many collaborations are prebuilt, the data representations they expect are also prebuilt. This prebuilt definition of a customer for a collaboration could be thought of as the "hub-data representation" of a customer.

Prebuilt collaborations are application independent, so they don't reference the specific fields of an SAP customer or a Siebel customer. Instead, this application-specific data is converted to the generic data expected by the collaboration. Business objects used by collaborations are called Generic Business Objects (also referred to as GBOs). The business objects initiating the flow of an event from an application are called Application Specific Business Objects or ASBOs (see Figure 4).

This is a one-to-many relationship where, for example, there is one Customer GBO and a potentially unlimited number of Customer ASBOs. As you can imagine, each application, from SAP to Siebel to PeopleSoft, has its own specific representation of a customer.

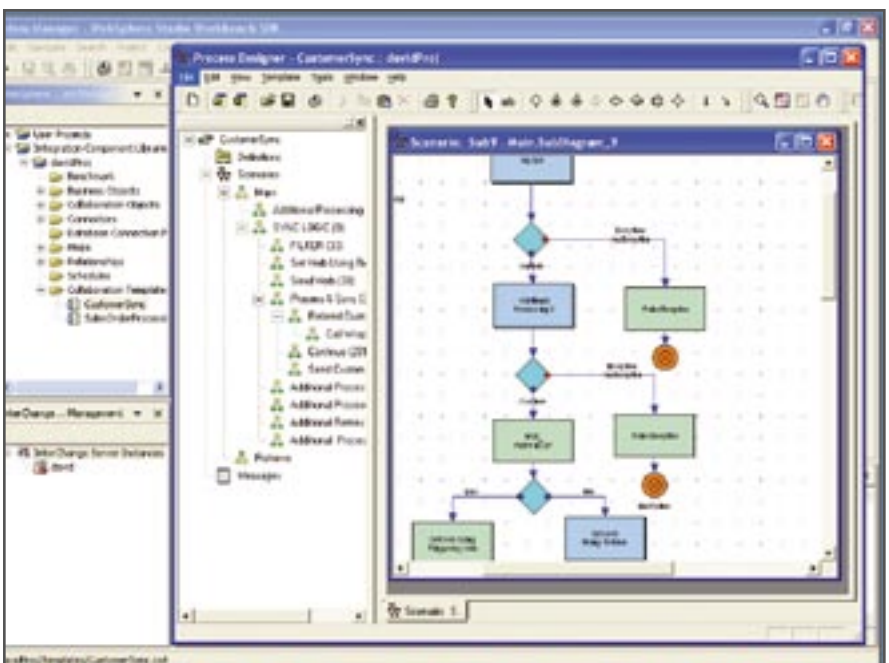

FIG 2: COLLABORATIONS CONTAINING BUSINESS LOGIC

WebSphere BI Server: Maps

Some of the WebSphere BI Server major components have been introduced. There is the hub (InterChange Server or ICS), the business process modules (Collaborations), the hub data (GBOs), and the application-specific data for the spoke applications (ASBOs).

Because the application-independent collaborations

Cameron Majidi has been a veteran consultant to Prolifics for nearly 10 years. As a senior enterprise solutions architect, he is part of Prolifics' team of WebSphere consultants that IBM calls on to service its most challenging customer requirements. Cameron has designed and implemented WebSphere solutions for customers around the world. He also has over five years of product design and management experience in application development and connector tools.
cmaj@prolifics.com

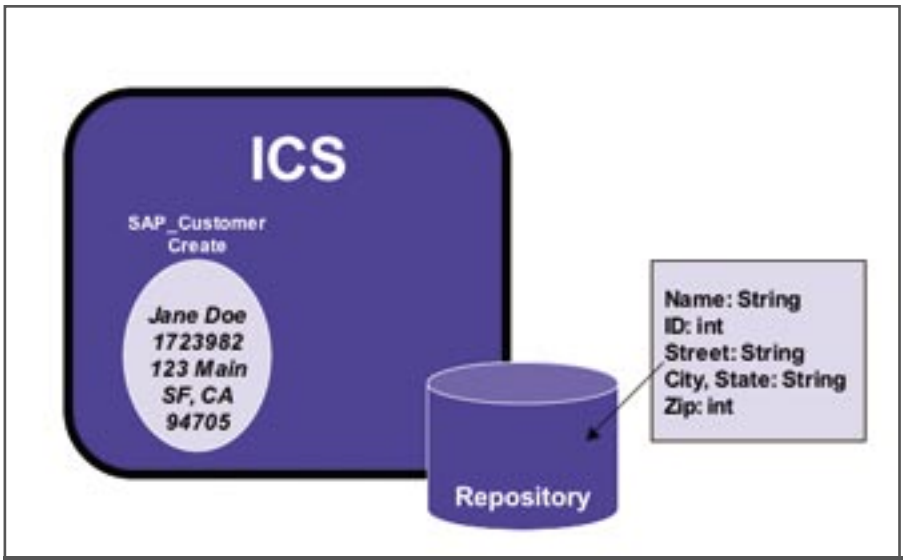

FIG 3: BUSINESS OBJECTS AND BUSINESS OBJECT DEFINITIONS

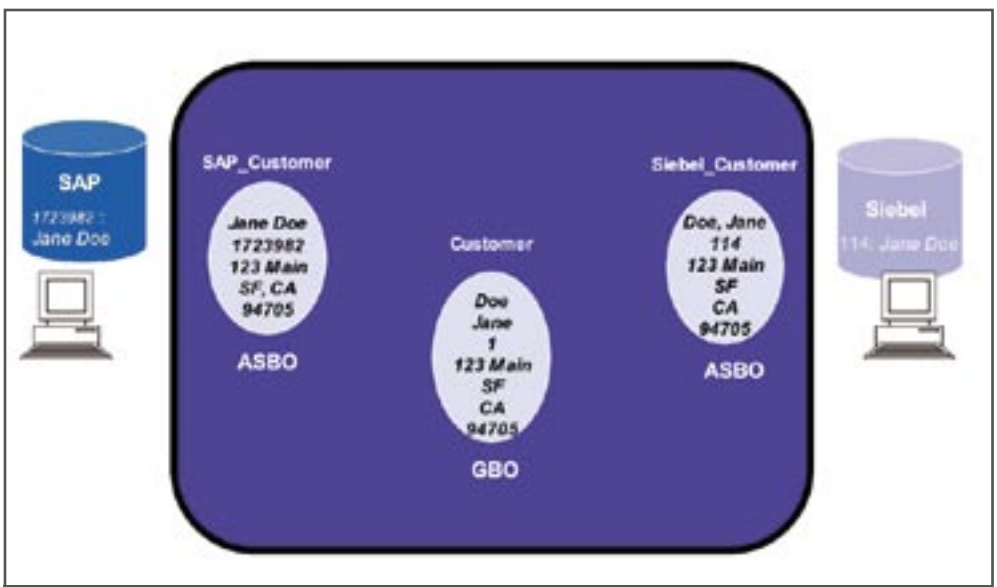

FIG 4: APPLICATION SPECIFIC BUSINESS OBJECTS AND GENERIC BUSINESS OBJECT

expect generic business objects, and the applications themselves provide application-specific business objects, there must be a conversion between ASBOs and GBOs. In WebSphere BI Server, this conversion of data is done with maps. The mapping objects are passed an ASBO and return a GBO if the data is heading from an application into a collaboration. Or, the mapping objects are passed a GBO and return an ASBO, if the data is heading from a collaboration into an application. The maps operate within the ICS hub, not as part of the spoke connectivity, for reasons explained later in this article (see Figure 5).

More on Business Objects

Recall that there is a GBO for each major entity expected by one of the prebuilt collaborations. For example, there is a Customer GBO (required by the Customer Synchronization collaboration) and a Sales Order GBO (required by the Sales Order Processing collaboration).

Actually, a Sales Order typically consists of one of more Items, so there is also an Item GBO. Similarly, a Customer typically has an Address, and an Address typically has one or more PhoneNumbers. Each separate entity has its own business object definition. So, there is a Sale GBO, an Item GBO, a Customer GBO, an Address GBO, and a PhoneInfo GBO.

Whether or not an application entity is defined as its own BO is often a decision based on reuse. Consider, for example, a phone number. Why does a phone number need a separate business object definition as compared to simply using a String representation? Although all phone numbers in the U.S. are 10 digits in length, we often have the additional 3-, 4-, or 5-digit extensions, because phone numbers outside the U.S. frequently vary in length. Also, consider the different types of numbers we have, ranging from fax to home to office to cell. A PhoneInfo business object definition could define the type of phone number and allow it to be reused in many other business object definitions. Creating a separate business object definition allows for a list of values. The WebSphere BI Server can't handle an array of PhoneNumbers unless PhoneNumber is defined as its own Business Object. There's no limit on the depth of a business object definition hierarchy. Nesting BOs allow for the representation of complex application schemas and modularity of design.

Business Objects and Maps

Given that a Customer actually consists of a hierarchy of BOs, there might be many maps involved – whether horizontally mapping from a Customer ASBO to a Customer GBO, and a Customer GBO to a Customer ASBO, or creating submaps to define parent-child transformations.

The total number of maps needed for an integration effort depends on

What//: are your Web applications saying to your customers?

diagnosys

Are your Web applications open 24 hours? Get the tips and best practices you need to keep your apps running at peak efficiency 24 hours a day with WebApp News, the free newsletter from the performance experts at H&W.

Sign up today at:
www.hwcs.com/wsj10.asp

You might as well be closed for business if your Web applications aren't running at top speed 24 hours a day. With poor Web performance costing businesses $25 billion a year, you need to find problems, and you need to find them fast.

You need DiagnoSys.

DiagnoSys is the first intelligent performance management solution that analyzes your J2EE and .NET applications from end to end. DiagnoSys starts analyzing quickly, gathers data directly from all application resources, and proactively finds and helps fix the real cause of problems – before they cost your business big time.

So when it's absolutely essential that your Web applications say "Open 24 hours," trust DiagnoSys.

H&W
www.hwcs.com | 1.800.338.6692

Bridging the Enterprise Computing Gap For 25 Years

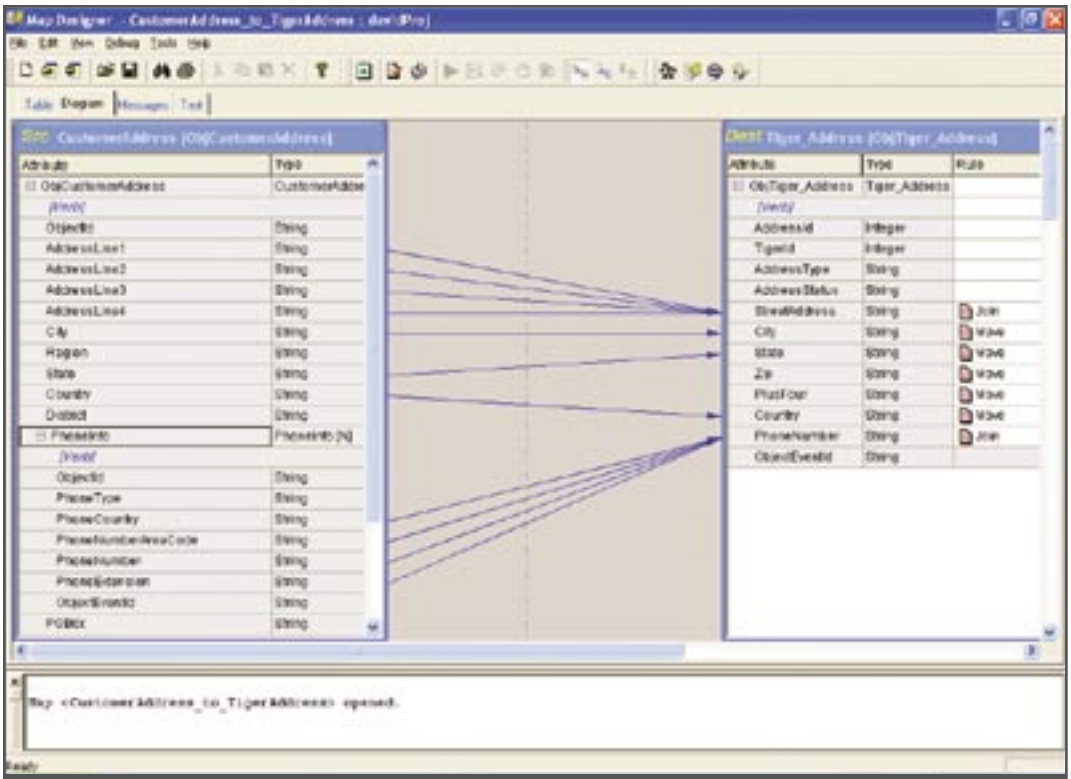

FIG 5: MAPS TO CONVERT ASBO AND GOB ADDRESSES

both the total number of applications involved and the total number of collaboration business logic modules.

WebSphere BI Server: Connectors

The final major component of WebSphere BI Server is responsible for extracting an ASBO from an application and sending it to the ICS, or receiving an ASBO from the ICS and inserting it into an application.

To be a spoke for the ICS, an application needs to be ICS-enabled. In other words, an application needs to be able to create ASBOs and send them into the ICS. The application also needs to be able to receive ASBOs from the ICS and take the appropriate action based on the verb. The component that connects an application to the ICS is called a Connector or an Adapter.

One task for a Connector is to determine when an application event (a change in its data) is something in which at least one of the collaborations is interested. It's the Connector's job to extract data from an application and insert data into an application. The Connector also connects the application to the ICS. Connectors can communicate with the ICS either through MQ (using the native MQ API or JMS) or IIOP.

As shown in Figure 6, the Connector is logically split into three parts:

- ***Connector Agent:*** the part that knows about the application (how to extract and insert data)
- ***Connector Controller:*** the part that knows about the ICS (specifically, about Maps and Collaborations)
- ***Connector Communication:*** the part in between, for communication between the previous two pieces: the Agent and the Controller

The SAP, Siebel, Clarify, and PeopleSoft Adapters are just four of the many prebuilt connectors available, called Application Adapters. If IBM doesn't already have a Connector for your application, you can write your own using their Connector Developers Kit. There are APIs for writing a connector in either C++ or Java. It's very rare, however, that anyone will have to build a custom Connector from scratch. In general, the preferred approach is to base a custom Connector on one of the prebuilt general-purpose Connectors. There are many Connectors that allow an ICS-enabling application, using popular technologies such as relational databases, queues, or files. These general-purpose Connectors are called Technology Adapters. Popular Technology Adapters include JDBC Adapter, JText (file) Adapter, MQ Adapter, and XML Adapter.

DATA FLOW

In this example, the desired integration business logic is Customer Synchronization from appA to appB. The prebuilt CustomerSync collaboration is used.

1. Connector Agent builds ASBO to represent Event
2. Connector Agent puts ASBO onto MQ
3. Connector Controller copies ASBO data from MQ and sends it to a Map
4. Map returns GBO to Connector Controller
5. Connector Controller calls appropriate Collaboration(s)
6. Collaborations perform business logic on GBO data
7. Business logic typically sends GBO to another application, a Connector Controller
8. Connector Controller sends the GBO to a Map
9. Map returns an ASBO to the Connector Controller
10. Connector Controller sends the ASBO to the Connector Agent (uses ORB for this)
11. Connector Agent performs the appropriate action based on ASBO

Figure 7 shows the Customer data flowing from appA to appB, but the final steps are missing. It's likely that the collaboration is waiting for a reply from appB. Otherwise, the collaboration doesn't know if the synchronization has been successful.

Updates to Data

The introduction to WebSphere BI Server's major components and the data flow description are preparation for a brief explanation of an advanced Enterprise Application Integration topic: updating data.

Conducting an update on customer data is very similar to creating a new customer until the final step of updating the data in appB. appB needs to know which customer to update. Applications typically generate their own internal representation of a unique key for every entry into a database table. Thus, Jane Doe on appA will have a unique key that isn't relevant in appB. It's not sufficient to tell appB to update Jane Doe because appB might have many different customers named Jane Doe. It's also unacceptable to tell appB to update the customer based on the unique key stored in appA.

The theme of hub-and-spoke once again drives the design of the mechanism for storing the cross-referencing information. Because there are many more than the two applications appA and appB, their unique keys are not stored in a single table. Instead, each key is stored in a table along with a hub-key, or a generic-key. The basic idea is that you never cross-reference appA directly with appB, or appC, or any other app. You cross-reference appA with the hub. In other words, instruct the hub to represent an incoming ASBO not by an arbitrary GBO, but by a particular GBO, with a unique identifier. When that GBO needs to be converted into a different ASBO for another application, a similar look-up is performed, mapping the GBO's unique ID to the unique identifier for the ASBO for appB to which it corresponds (see Figure 8).

This clarifies an earlier claim about the mapping being done in the hub instead of in the spokes. The cross-referencing tables are stored in the repository instead of on the separate applications. Also, the cross-referencing is done as part of the mapping. So, the mapping must be handled in the hub.

Beyond the Firewall: Java to Application

Everything discussed so far assumes that a company runs its business processes using data exclusively from its internal software applications. In this case, a Connector gets data from the application into the ICS in preparation for a collaboration to automate an integration business process.

A second option is to call a collaboration directly from another application, which would bypass many steps of the typical data flow outlined previously. A direct call to a collaboration would mean that there would be no Connector Controller, and therefore, no process in charge of calling a map to translate ASBOs to GBOs and GBOs to ASBOs. Given that the collaboration is expecting a GBO, there are two options:

1. Pass a GBO directly to a collaboration. Because you're likely to be in control of the code for whatever application you're using to call a collaboration, generate data that matches the GBO definition required by the collaboration.
2. Configure the collaboration so that it will call a map. This is the more consistent option of the two, allowing the collaboration to change the GBO data specification without affecting the call from the external application. The external application can pass ASBO data, knowing that an ASBO to GBO map will handle the specifics of preparing the data for the collaboration. Similarly, from the ICS back to the application, a GBO to ASBO map is used.

Note that the code that calls the collaboration could be a Java Servlet or an EJB. This allows Java applications, such as those running in an application server, like WebSphere Application Server, to integrate very closely with collaborations in ICS. The API used for accessing a collaboration directly is called the Access Framework. This mechanism uses IIOP to communicate with the server and is a tighter (synchronous) style of interaction with the server than connecting an application with ICS using the MQ Adapter and JMS.

Because the ICS server can act as a general gateway to all the non-Java legacy applications in a given enterprise's environment, using the Access

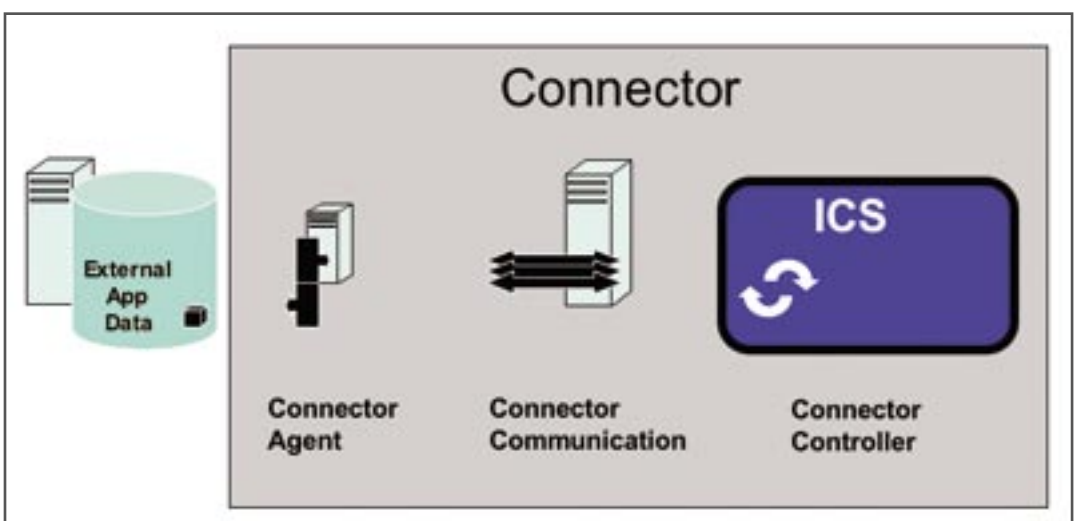

FIG 6: THE CONNECTOR IS COMPRISED OF THREE PARTS

Framework to call collaborations in the ICS is a good way to integrate a J2EE application with any number of legacy back-end systems. Each such call would be the same – amounting to sending a BO to the ICS server. As a result, the Java code in the J2EE application wouldn't have to implement different mechanisms for connecting to, say, SAP, PeopleSoft or a CICS application. In each case, it would use the Access Framework as a generic RPC mechanism, and let the ICS and its associated Connectors handle all of the details regarding which back-end systems to connect to and how.

Beyond the Firewall: Trading Partner Data to Application

So far, we have discussed how WebSphere BI Server can integrate and act upon data from software applications, both in-house and across firewalls, as well as data from other sources, such as Web browsers. Today's companies, however, handle data from various other locations such as e-business exchanges with trading partners and marketplaces. IBM offers a product for managing trading partners called WebSphere BI Connect, which is complementary to WebSphere BI Server. WebSphere BI Server is the integration and business process hub, whereas WBI Connect is the B2B trading partner management and routing hub.

Using WebSphere BI Connect, a company can easily manage and securely communicate with hundreds of trading partners. Using both WebSphere BI Connect and WebSphere BI Server, the data flowing to and from trading partners can also flow through collaborations. If a company's business processes are already being implemented by collaborations running in the ICS, the mechanism is in place and ready to handle external data as well.

WebSphere BI Connect is designed to handle receiving and sending EDI, XML, and various documents and secure transmission with digital certificates, encryption, and decryption. However, WebSphere BI Connect doesn't process the content; it's the job of WebSphere BI Server to process the content from a variety of sources and in a variety of formats.

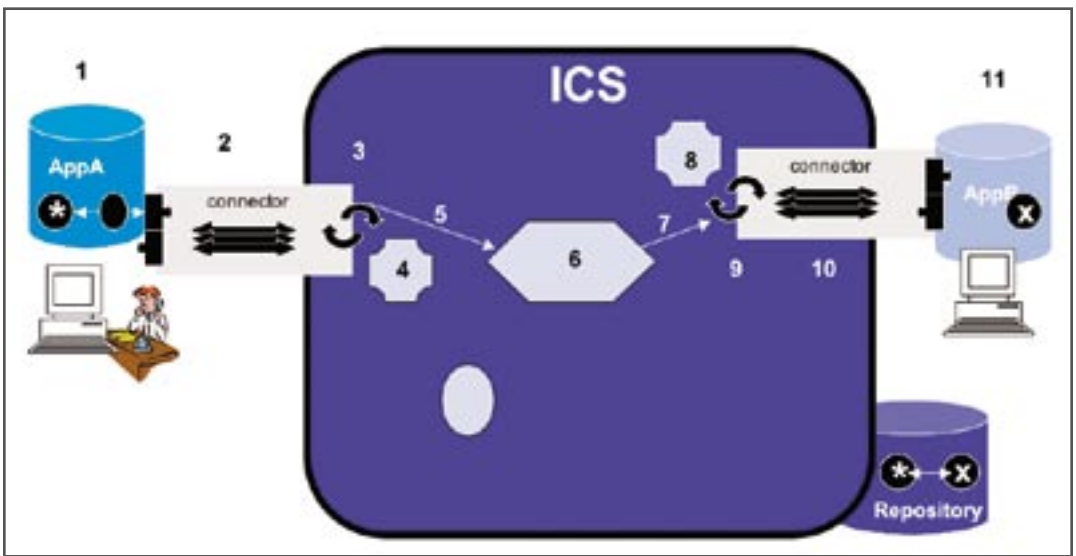

FIG 7: EXAMPLE DATA FLOW OF CUSTOMER SYNCHRONIZATION BETWEEN APP A AND APP B.

Exposing Collaborations as Web Services

Finally, the discussion of using WebSphere BI Server and WebSphere BI Connect for e-Business wouldn't be complete without mentioning that collaborations are J2EE objects and can be exposed as Web services. Therefore, if a company desires, it can design collaborations for both internal and external use.

Conclusion

There are many more details about IBM's WebSphere BI Server that are beyond the scope of this article. However, we have touched on all of the major components and armed you with both terminology and a base understanding from which to begin further WebSphere BI Server education and implementation of integration business processes.

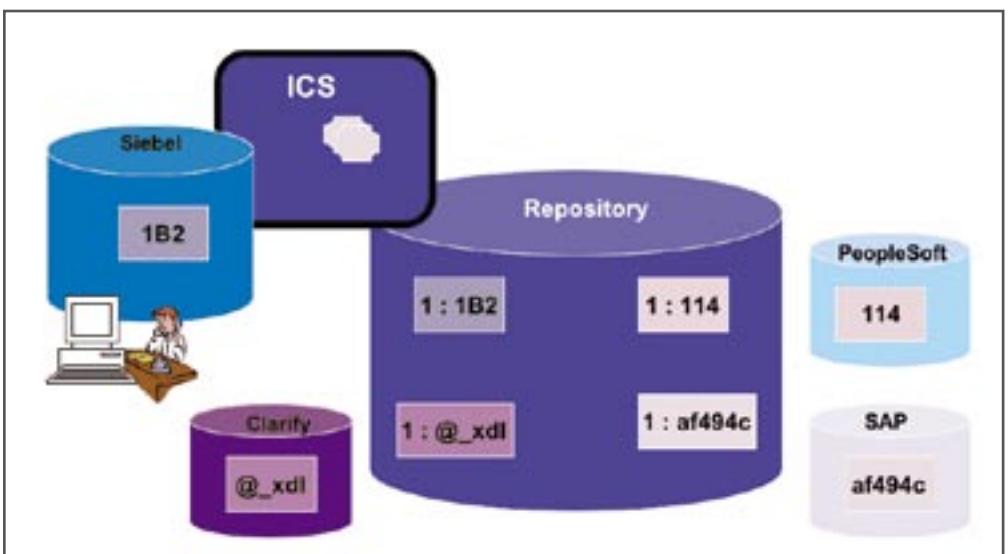

FIG 8: CROSS-REFERENCING OF UNIQUE KEYS USING A HUB-AND-SPOKE DESIGN

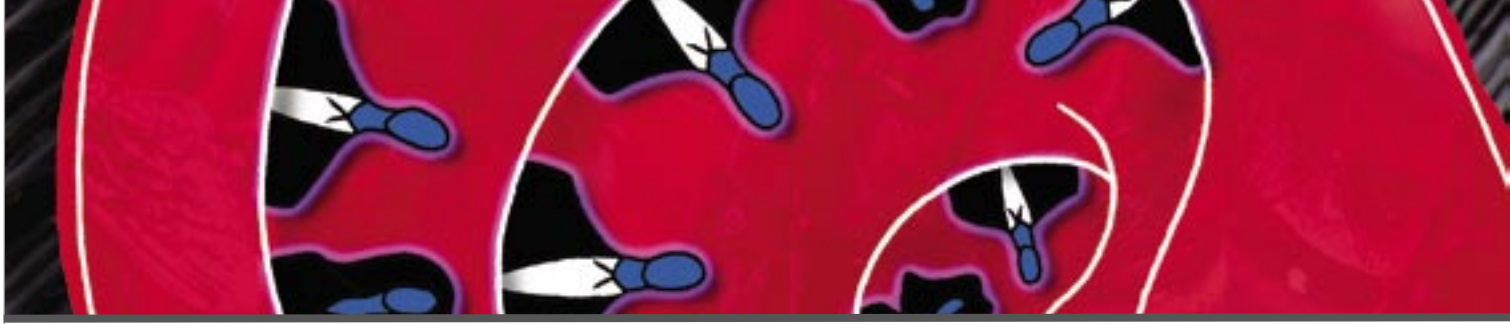

Understanding JavaServer Faces

Assemble your Web applications effectively

BY JEFFREY LIU & HELEN WONG

Jeffrey Liu is a software developer on the IBM Rational Studio Web Services Tools Team at the IBM Toronto Lab. He delivers tools that ease the Web services development effort. He is also an active contributor to various WebSphere domains and magazines. jeffliu@ca.ibm.com

Helen Wong is a software developer on the C/C++/Fortran Compiler Team at the IBM Toronto Lab. She is an active contributor to IBM developerWorks. helenhyw@ca.ibm.com

The Internet has no doubt become the most popular source of information today. More and more businesses have adopted the Internet as the vehicle to display their business information, advertise their products, and maintain customer interactions. The activity of developing Web applications occupies a fair amount of resources, time, and effort. Doing so in an efficient and timely manner directly contributes to a business's success.

In the past, technologies such as JavaServer Pages (JSP) have provided a primitive approach to develop HTML-based Web applications using Java. However, they have failed to provide a standard and reusable UI framework. The JavaServer Faces (JSF) technology offers a standard and effective way to manage and develop reusable UI components, which together assemble a Web application. This article examines JSF technology in both the business and technical horizons. You will be presented with the benefits of JSF and a scenario demonstrating the JSF support in the IBM WebSphere Studio Application Developer product.

JSF: A Business Perspective

The Internet formulates a virtual environment for companies to conduct their business. More and more companies have exploited the Internet as a means to gain a competitive edge. Through the use of the Internet, companies can increase visibility, increase accessibility, and at the same time, reduce cost, which allows for greater profit and success. In this digital marketplace, a company's Web site represents its value and serves as the starting point for customers to learn about the company. Companies spend thousands and thousands of dollars to research, design, and implement their Web sites, using them to advertise their products, carry out their day-to-day businesses, and service their customers.

As a company grows, its Web site must grow along with it. Specifically, a company's Web site must reflect the latest development of the company. This can be an announcement of a new product release or the addition of a new service. The former may require as little effort as adding a few Web pages and links to the Web site, while the latter may require substantial effort by adapting to the business logic that the new service offers.

The ability to rapidly develop a Web site to support business changes poses a challenge to many companies. Consequently, a reusable, scalable, and extensible user interface (UI) framework for developing Web applications is undoubtedly desirable. In the Java domain, such a framework did not exist until the introduction of JSF, which serves as the strategic solution for developing Java Web applications.

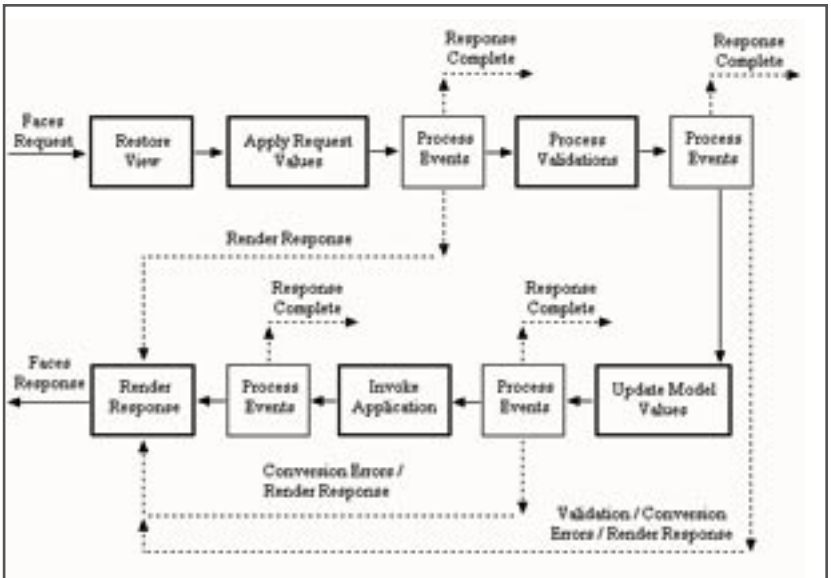

FIG 1: FACES REQUEST PROCESSING LIFE-CYCLE

In today's customer-oriented economy, even the best technology does not guarantee market acceptance. Factors such as benefits to business, cost of adoption, and availability of service and support from tools influence a business's decision on whether or not to adopt a new technology. Businesses must understand the implications of a new technology on each of the above concerns.

JSF provides a UI framework that eases the process of designing, developing, and maintaining Java Web applications. JSF brings out the notion of roles in Web development, clearly defining a separation between business logic and presentation logic. This allows the UI (i.e., Web pages) to be developed independently of the business logic, enabling a loose coupling between presentation and business implementation. This, in turn, improves reusability. Previous technologies such as JSP and servlets fall short of reusability because of the mixed usage of UI-centric language HTML and the application-centric Java language in the source files. JSF eliminates this problem by defining different roles in Web development. These roles are page author, component writer, application developer, tool provider, and JSF implementor.

Page authors, who are fluent in markup and scripting languages such as HTML and JavaScript, specialize in the creation of the user interface. They focus on the visual design of the Web application, and are responsible for assembling the user interface using components provided by component writers.

Component writers create reusable UI components, which are Java classes that represent UI controls, such as an input text field on a HTML page. Component writers are also responsible for supplying a render that transforms a UI component into a format that the clients expect, such as HTML.

Application developers provide the underlying business logic. This can be an Enterprise JavaBean (EJB) component that exposes the application functionalities. Application developers are proficient in programming languages such as Java.

Tool providers and JSF implementors are not directly involved in the development of the Web applications. Typically, they are vendors, such as IBM, that provide the development tools to assist in other roles in developing JSF-based Web applications, and provide the platform to run those applications.

Unlike the JSP and servlet technology, such a role-based architecture does not require developers to be skilled in both markup/scripting languages and application programming languages. This architecture allows developers to specialize, to have a field of expertise, which results in greater productivity, better code maintenance, and shorter development cycles. With JSF, businesses may also be concerned about the effort to port existing code such as JSP to the JSF framework. In fact, existing JSP components can coexist with any JSF-based Web pages within the same Web application.

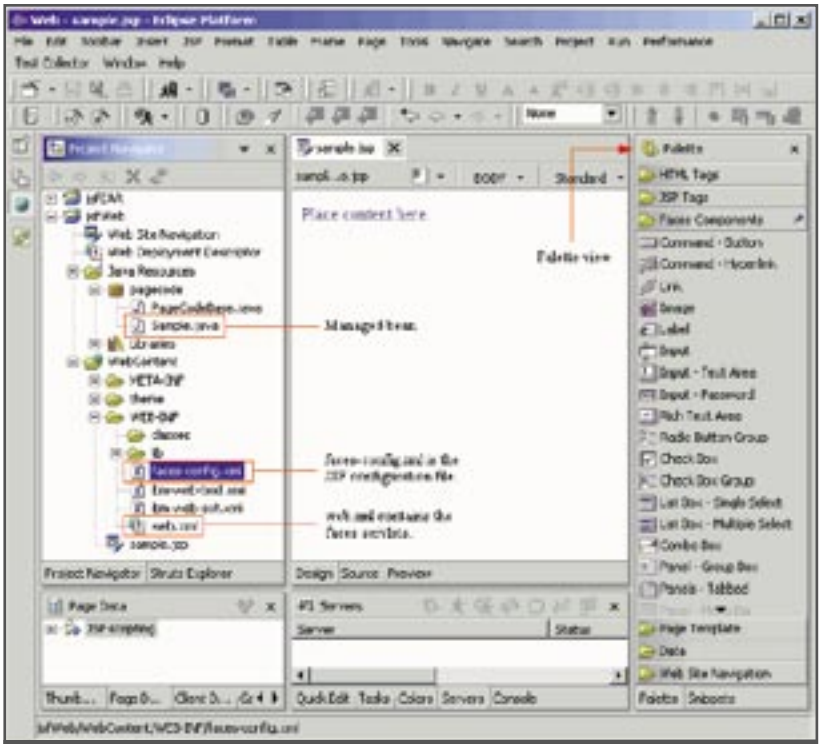

FIG 2: WEB PERSPECTIVE

To make JSF more intuitive and appealing to potential users, tools that support the JSF technology play an important role in assisting novice users. Because the majority of the learning effort lies within the JSF syntax, the business incentive to adopt JSF blossoms if development tools can hide the complex details of JSF from the users, offering easy-to-use features such as drag-and-drop action for UI components, lowering the barrier of entry.

JSF: A Technical Perspective

Web development on the Java 2 Enterprise Edition (J2EE) platform has evolved from a primitive approach – servlet technology – to a more architectural approach

"More and more companies have exploited the Internet as a means to gain a competitive edge"

"The WebSphere Studio Application Developer version 5.1.2 marks a significant milestone in developing Web applications on the J2EE platform"

such as that used in JSP components and the JavaServer Tag Library (JSTL). This transition has dramatically changed the development environment from one that fuses Java and HTML together to one that provides better separation between writing UI and application code. Things are moving in the right direction, but they have yet to fulfill the increasing business demands on Java Web development. Even when using JSP and JSTL technology, Web developers are often required to create their own frameworks to support their Web applications. JSF recognizes the lack of a reusable UI framework with current Java technologies, and therefore offers a standard to ease such effort.

JSF employs an entirely new concept of building Web applications in the Java domain. The design of JSF utilizes a meta-model, which is a collection of UI components, to continuously represent the current state of the client UI. This is quite the opposite of the way previous technologies worked, in which a thin-client approach was used to maintain a stateless client UI. The use of a meta-model grants JSF-based Web applications the same capabilities of state management, event handling, and validation that thick-clients such as a Java Swing-based applications enjoy.

As discussed earlier, in the world of JSF, presentation logic and application logic can be developed independently of each other. However, a Web application cannot be deemed complete until these two pieces are combined. This involves the binding of application data, such as a JavaBean field, to UI components. The synchronization between application data and the client UI is handled by the JSF engine provided by the JSF implementor. In addition to data binding, JSF also allows other types of connections, for example, an association of a JavaBean method with a UI action. These connections are described using XML markups in faces JSP files, making these files free of scriptlets, which in turn increases readability. In Figure 1, we illustrate the faces request processing life-cycle.

Similar to other HyperText Transfer Protocol (HTTP)-based Web applications, the faces request processing life-cycle begins when a client makes an HTTP request via a Web browser. Such requests can be categorized into two types: non-faces request and faces request. A non-faces request originates from a non-UI component. For example, the client enters a Uniform Resource Locator (URL) into the browser and loads a faces JSP file. A faces request originates from a UI component such as a faces submit button. In a typical faces request scenario, the life-cycle goes through six different phases:

- ***Restore view:*** Upon receiving a faces request, the JSF engine retrieves the faces context instance associated with this user session. The faces context instance contains a view (collection) of UI components that represents the client-user interface.
- ***Apply request values:*** In this phase, the states of the UI components are synchronized with those on the client-user interface. This is done by reading the appropriate data from the HTTP request, such as headers and parameters. As part of the synchronization, events may be triggered, and the registered event handlers will be invoked. This enables

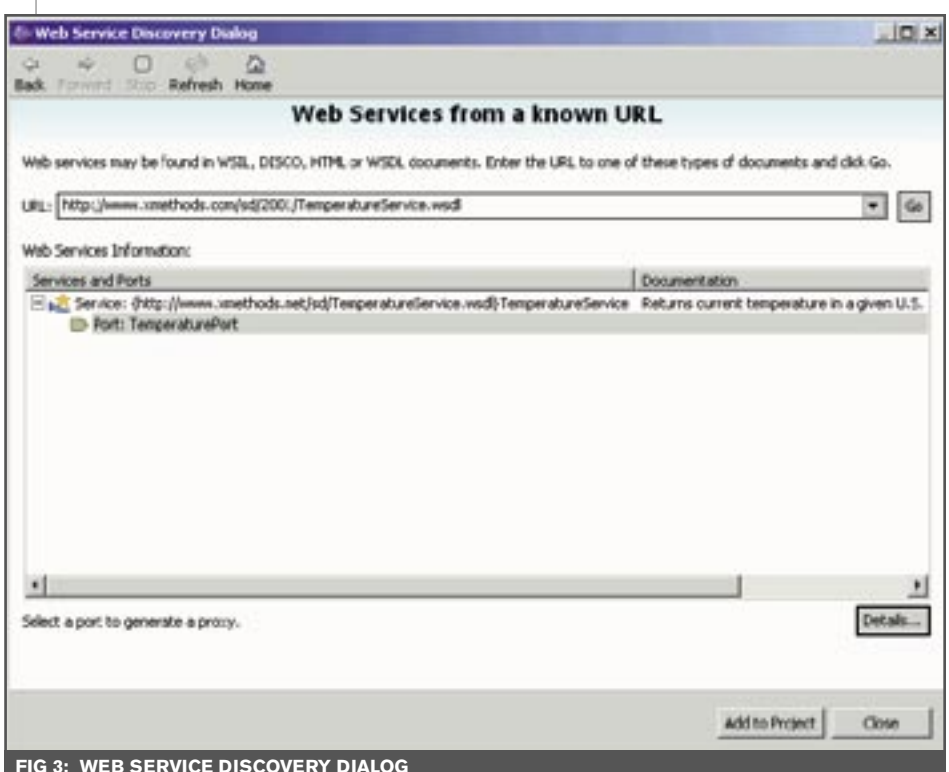

FIG 3: WEB SERVICE DISCOVERY DIALOG

Build WebSphere Portlets FAST*

**and change them even faster.*

Out-of-the-box portlets aren't enough.
But, you already know that.

Driving directions, stock quotes, and news are nice, but they're not what you need to run your business. What you need are customizable portlets that bring your enterprise IT assets to your enterprise portals.

Bowstreet™ Portlet Factory is the first product to automate the rapid development, deployment, and maintenance of highly adaptive portlets that plug right into the IBM WebSphere Portal stack.

BOWSTREET PORTLET FACTORY
TRY IT FREE FOR 60 DAYS.
WWW.BOWSTREET.COM/EASY

bowstreet

J2EE IS HARD. WE MAKE IT EASY.

server-side event handling for client-generated events. Such an event model is something that previous technologies did not possess.

- ***Process validation:*** After setting the states of the UI components, validator instances registered with each UI component are triggered. Errors are immediately reported back to the client.
- ***Update model values:*** The purpose of this phase is to update the application data, changing the value of a JavaBean field, for example.
- ***Invoke application:*** The invoke application phase processes the action submitted by the request, such as invoking the JavaBean method associated with the submitted action.
- ***Render response:*** This is the final phase in the faces request processing life-cycle. This phase transforms the UI components into a format that the client expects, such as an HTML document.

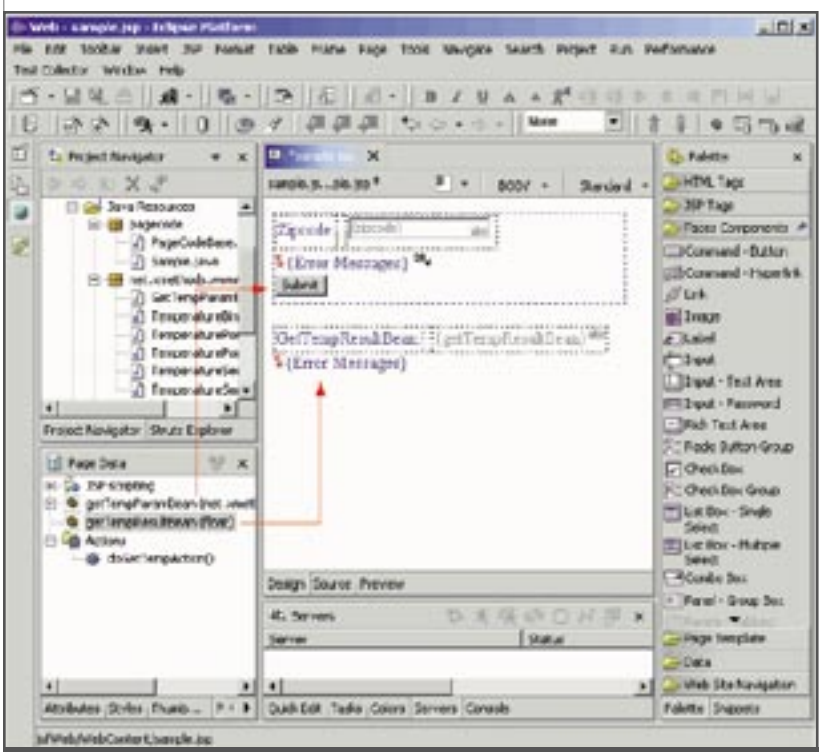

FIG 4: DRAG AND DROP A WEB SERVICE

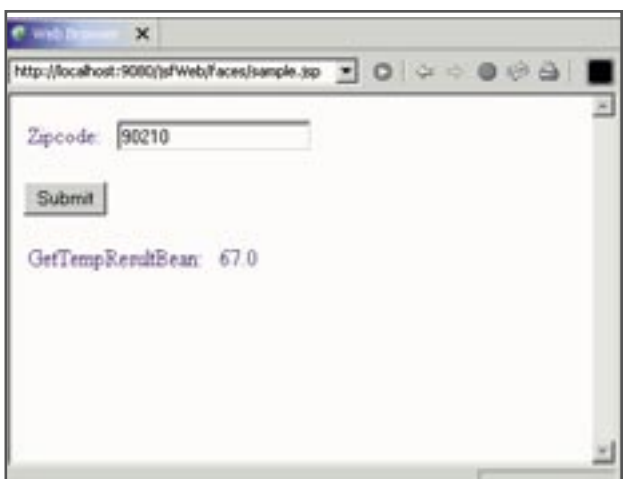

FIG 5: RESULT OF INVOKING THE TEMPERATURE WEB SERVICE

"With the right tools, developing a JSF-based Web application has never been easier"

JSF Support in IBM WebSphere Studio Application Developer

The WebSphere Studio Application Developer version 5.1.2 marks a significant milestone in developing Web applications on the J2EE platform. It not only provides support for the JSF specification version 1.0, but it also allows seamless integration between JSF and other tools.

Application Developer provides a wizard to guide the user through the creation of a template faces JSP file. This wizard automatically generates the essential faces artifacts, which include the faces servlets in Web.xml, the JSF configuration file faces-config.xml, and a managed bean. The javax.faces.webapp.FacesServlet faces servlet listens for incoming requests and initiates the faces request processing life-cycle. The faces-config.xml file defines configuration settings such as the managed bean facility. The managed bean facility contains one or more managed beans, which serve as the integration points between the presentation logic and the business logic. This wizard serves as a starting point for developing a JSF-based Web application. Other JSF-related tools can be found in the Web perspective, as shown in Figure 2.

From the palette view, users can drag and drop various UI components onto the faces JSP file. Upon dropping a UI component, the XML markup in the faces JSP file and the corresponding Java code for the UI component in the managed bean are automatically generated. This shields the users from the complicated details of JSF.

To leverage the potential of JSF, WebSphere Studio Application Developer provides support for more than just the standard set of UI components. It also makes available other custom UI components, as well as integration with other tools, such as Web services. The Web Service Discovery Dialog demonstrates such integration by allowing users to discover a Web service, generate a JavaBean client for that Web service, and drag and drop the corresponding Web service operations onto a faces JSP file. Figure 3 shows the discovery of a temperature Web service using the Web Service Discovery Dialog.

Upon discovering a Web service,the generation of a JavaBean client is just one click away. This temperature Web service has a single operation that takes in a getTempParamBean as input and returns a getTempResult-

Bean as output. As shown in Figure 4, these two parameters, which can be found in the page data view, have been dragged and dropped onto a faces JSP file. In addition, a submit button, which is bound to the Web service operation, is also generated. When a user clicks on this button, a faces request containing the value of the getTempParamBean is submitted to the faces servlet.

Now, let's uncover the underlying details of the tool. As part of dragging and dropping a parameter, the tool will generate the XML markup in the faces JSP file and the UI component in the managed bean. Listing 1 shows the XML markup for the getTempParamBean parameter, Listing 2 shows its corresponding UI component, and Listing 3 shows the associated application data.

The value attribute of the <h:inputText> element in Listing 1 defines the value binding between the zip code field of the getTempParamBean and the HtmlInputText UI component. Besides the value binding, there's also an action binding shown in Listing 1. The action attribute of the <h:commandExButton> element defines the binding between the doGetTempAction method and the HtmlCommandExButton UI component. During the faces request processing life-cycle, specifically in the apply request values phase, the value of the HtmlInputText UI component is synchronized with the value submitted in the HTTP faces request. Next, in the update model value phase, the zip code field of the getTempParamBean is updated with the value of the HtmlInputText UI component. In the invoke application phase, the submitted action is then processed. The doGetTempAction method is called, in which the temperature Web service is invoked. Figure 5 shows the result of using this faces JSP file to invoke the temperature Web service.

Conclusion

As demonstrated, to manually develop and assemble such applications without the aid of any tools could be very time-consuming and error-prone. With the right tools, developing a JSF-based Web application has never been easier.

The innovative JavaServer Faces technology is a big step forward in Web development on the J2EE platform. With this new standard, businesses benefit from a comprehensive, reusable, and extensible UI framework. The JSF design offers excellent state management, event handling, and validation solutions that modern applications can utilize. As demonstrated in this article, using IBM WebSphere Application Developer to develop a simple JSF-based Web application to invoke a Web service can be done in just a few steps. One doesn't have to be both a Java guru and a Web expert to take advantage of JSF.

References

- *JavaServer Faces specification:* www.jcp.org/en/jsr/detail?id=127
- *WebSphere Studio version 5.1.2 trial download*: www-306.ibm.com/software/websphere/info/easierjava/index.jsp

LISTING 1: XML MARKUP FOR GETTEMPPARAMBEAN

```
<h:inputText
  id="text1"
  value="#{pc_Sample.getTempParamBean.zipcode}"
  styleClass="inputText">
</h:inputText>
…
<hx:commandExButton
  id="button1"
  styleClass="commandExButton"
  type="submit"
  value="Submit"
  action="#{pc_Sample.doGetTempAction}">
</hx:commandExButton>
```

LISTING 2: UI COMPONENT FOR GETTEMPPARAMBEAN

```
protected GetTempParamBean getTempParamBean;
protected HtmlInputText text1;
…

public GetTempParamBean getGetTempParamBean() {
  if (getTempParamBean == null) {
    getTempParamBean = new GetTempParamBean();
  }
  return getTempParamBean;
}

public String doGetTempAction() {
  if (temperaturePortTypeProxy == null) {
    temperaturePortTypeProxy =
      new TemperaturePortTypeProxy();
  }
  try {
    getTempResultBean =
      temperaturePortTypeProxy.getTemp(
        getGetTempParamBean().getZipcode());
  } catch (RemoteException e) {
    logException(e);
  }
  return "";
}

protected HtmlInputText getText1() {
  if (text1 == null) {
    text1 =
      (HtmlInputText)findComponentInRoot("text1");
  }
  return text1;
}
```

LISTING 3: APPLICATION DATA FOR GETTEMPPARAMBEAN

```
public class GetTempParamBean {
  public String zipcode;

  public String getZipcode() {
    return zipcode;
  }

  public void setZipcode(String zipcode) {
    if (zipcode != null) {
      this.zipcode = zipcode;
    }
  }
}
```

Calming the configuration beast

WebSphere Migration

BY RAMAN SUD

Raman Sud is the vice president of engineering for mValent, developer of the Infrastructure Automation Suite. Sud has 20 years of experience delivering mission-critical software for enterprises and telecommunication service providers leveraging distributed development and building integrated teams in the U.S. and India.

In the coming months, hundreds of corporations will migrate thousands of applications from the WebSphere 4.x platform to WebSphere 5.0. The reasons for doing so are both many and compelling, not the least of which is that support for the WebSphere 4.x platform will cease to exist in a matter of months.

However, beyond that practical reason, WebSphere 5.0 provides a giant leap in flexibility, functionality, scalability, and expandability with its new distributed architecture. This makes it well worth their while for most companies to upgrade. However, upgrading from 4.x to 5.0 has its difficulties. None is more daunting than the challenge of porting application configurations from the WebSphere 4.x to the distributed WebSphere 5.0 architecture.

For every application deployed on the WebSphere platform, IT personnel must carefully enter dozens of pieces of configuration information to ensure that the application works properly. They also must engage in similar configuration activities across all "dependent" components in the application architecture, such as database servers, load balancers, and firewalls. One slip of the finger can cause a bad configuration, which results in poor application performance or, worse yet, application service outages.

WebSphere 5.0's new distributed architecture means that its configuration requirements are completely different from those of previous versions of the platform. Therefore, if one were to take a WebSphere 4.x application and deploy it on WebSphere 5.0, the application simply would not work. IT personnel have to reconfigure the application so that it works properly on the new platform. However, this can be a daunting task because a simple configuration file in WebSphere 4.0 must be broken apart, extended, and distributed across 30 separate configuration files in WebSphere 5.0.

IBM has published an excellent book entitled *Migrating to Websphere v5.0 – an End-to-End Migration Guide*. While this provides step-by-step instructions on how to migrate configurations, IBM does not provide any tools for automating this process. As a result, IT departments are forced to do the migrations manually, one application at a time, and must hire additional staff simply to handle the increased complexity in configuration management.

The manual configuration process is the reason why it typically takes roughly half a day to set up a single application server. For a large company that may need to implement hundreds or even thousands of instances of WebSphere, this process represents an enormous productivity drag. Additionally, because each configuration requires keen attention to detail across a multistep process, it is highly error-prone, leading to deployment delays and application downtime.

Part of the problem is that the WebSphere administrative console is built to support a "one-at-a-time" manual process, not a "one-to-many" automated one, in which entire nodes, cells, or hives can be configured and managed collectively. And, there is no reusability of configurations. With every application deployment, IT personnel must "reinvent the wheel" and start from scratch with configurations.

This leads to inconsistent configurations, even among application servers that are supposed to be configured identically. Furthermore, there is no shared knowledgebase for these configurations, because each machine is configured "in a vacuum" by a specific individual. If he or she leaves the department, other staffers have to determine why that person configured the machine in a particular way.

Many companies today are overcoming this problem by adopting new automated infrastructure configuration management solutions. These products have become available to the market throughout the last year, creating a layer of abstraction between IT staff and infrastructure, so they no longer have to understand the specific configura-

tion requirements of each infrastructure component.

These solutions provide a configuration management database so IT departments can create a centralized "library" of configurations that can be examined and reused when appropriate. This reusability can reduce the time of deployment for a new application server from half a day to less than half an hour, radically improving the overall productivity of the IT department. It also eliminates manual configuration errors, enabling "service-oriented IT," in which applications can be flawlessly configured, deployed, and modified on a moment's notice, to keep pace with constantly changing business requirements.

The following sections go into greater detail as to how this type of solution can not only ease the migration to WebSphere 5.0, but also provide an ongoing, consolidated solution for managing configurations across the entire application lifecycle. For the purposes of illustration, we will use a real-life example of a WebSphere migration at a large publisher, which uses an automated configuration management solution.

2004: A Migration Odyssey

A few months ago, the IT department in one of the world's foremost publishers of business information was faced with two substantial challenges. First, the department needed to move all of its production configuration data from a relational database to XML. Second, the company needed to migrate from WebSphere 4.x to WebSphere 5.0, which meant creating and managing all new application configurations for the 1,500 application servers deployed across the company's test, certification, production, and disaster recovery environments.

Both of these challenges required a similar technical solution: a consolidated way in which to create, manage, and deploy application infrastructure configurations. With no commercial solutions geared specifically toward managing these configurations across a large-scale, distributed environment, the department had developed homegrown tools and scripts for managing WebSphere 4.x configurations. However, because WebSphere 5.0 had a completely different management interface and configuration management model, those tools and scripts were obsolete.

Because it was not practical to create new tools to manage WebSphere 5.0, it became clear that the only way to meet these technology challenges was to use an age-old solution: brute force. The company's IT department would simply have to undertake an incredibly time-consuming and laborious process of manually creating and managing configurations for all 1,500 application deployments on the WebSphere platform.

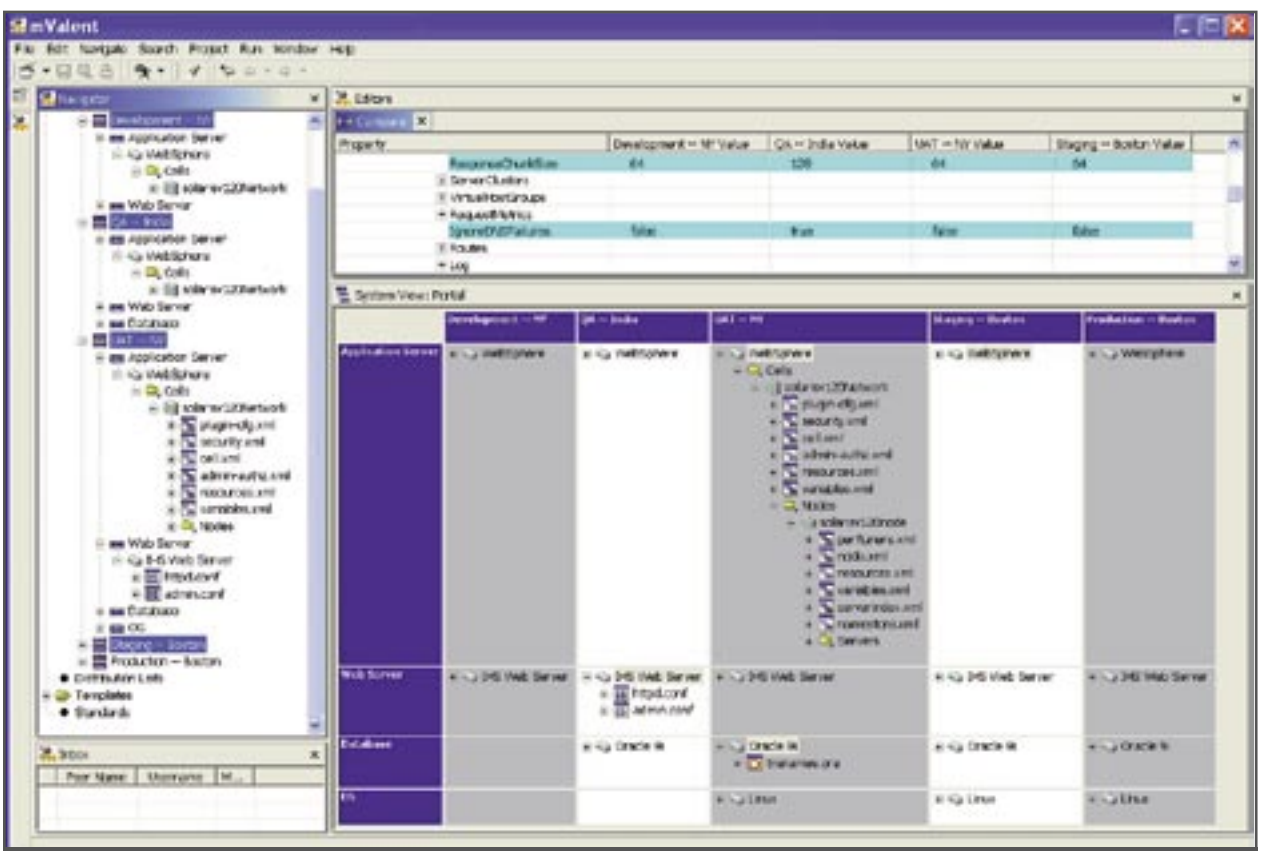

FIG 1: MANUAL CONFIGURATION SOLUTION

But just when it appeared that thousands of valuable IT hours would be consumed by manual labor, one of the company's IT executives was introduced to a configuration management solution that automates the creation, management, and maintenance of IT infrastructure configurations. This solution enables enterprises to automatically build, manage, and maintain complex infrastructure configurations throughout the entire application life-cycle, across distributed and off-shore locations. By adopting this solution, the IT department would be able to eliminate the manual processes involved with configuration creation and management, and instead automate them through a single, scalable enterprise solution.

Curing Configuration Chaos

Using manual techniques, the company estimated that for every WebSphere-based application migration from 4.x to 5.0, someone on the IT staff would have to create and manage more than 20 identical configurations across the company's test, certification, production, and disaster recovery environments. To put this into perspective, let's examine the cost savings of employing an automated configuration management solution vs. using manual techniques (see Table 1).

Table 1 illustrates the incredible amount of waste involved with manual configuration, and it does not even touch upon the thousands of opportunities for human error every time a configuration is manually implemented on the WebSphere 5.0 platform. By automating the management of these configurations, the labor savings alone justified the investment in the infrastructure configuration management solution. When factoring in the additional benefits of eliminating the inevitable human errors that arise with manual configuration management, the investment in the solution became a "no brainer."

> **"Configuration-related downtime has virtually disappeared"**

The company estimated that it would have required roughly two years of work to build a tool that could automate configuration management for the WebSphere 5.0 platform. By moving to a commercial solution, it was able to adopt the solution within a matter of weeks and slash migration time by up to 75% per application.

Ongoing Benefits

The benefits of automated configuration management were felt immediately in the form of radically reduced migration time and the elimination of manual configuration-based errors. The solution also enabled "one-to-many" configuration management, so configurations can now be updated and managed across all application elements from a single console.

This has enabled the IT department to keep all of its configurations consistent across systems. This had always been the department's goal, but when using manual, one-at-a-time configuration processes, it was impossible to accomplish, and, in reality, no two supposedly identical application servers were ever configured exactly alike.

Likewise, the solution automates the management of dependent configurations, such as those of Web servers and database servers. This enables the IT department to implement the configuration changes automatically – every time they need to make a configuration change on the WebSphere platform. Historically, this has been one of the most common areas of configuration errors, not only because it requires configuration changes across a number of different application elements, but also because those elements are usually managed by different people within the IT organization.

Diagnosis and repair of configuration problems has also been simplified. In the past, whenever there was a configuration-related problem with an application, the IT department had to manually compare the bad configuration parameters with a known good configuration to find the problem. It could take hours to

TASK DESCRIPTION	TYPICAL TIME USING MANUAL TECHNIQUES	AUTOMATED CONFIGURATION MANAGEMENT	SAVINGS
Build & deploy new App Server	4 hours	10 mins	96%
Replicate an existing application environment (Web, App, DB)	1 day	15 mins	98%
Promote a WebSphere environment from Dev, to QA, to Certification, etc.	2 days	15 mins	99%
Compare WebSphere related technology stack to determine differences	8 hours	15 mins	96%
Troubleshoot new 5.x environments	4 hours	15 mins	93%
Visualize Websphere environments	Days, manual effort using Visio type tool	Mins, automated	99%

TABLE 1: AN EXAMINATION OF AN AUTOMATED VS. MANUAL TECHNIQUES

find the bad configuration using this technique. This was more than just a productivity problem, however. It actually jeopardized the company's ability to generate revenue, as millions of end users rely on its applications for vital information and purchase millions of dollars of archived content every day. Every minute of application downtime represented lost revenue.

With its new automated infrastructure configuration management solution, the IT department can identify and repair configuration inconsistencies automatically, within a matter of seconds, ensuring that it never has revenue interruptions due to application outages.

Automated configuration management has also simplified the migration of configuration data from relational to XML. All of this data can be stored and managed in the central repository and pushed out to application components as needed.

Before adopting configuration management software, any change in an application required people to go into multiple systems to input manual changes. This led to frequent errors and inconsistencies, often resulting in application malfunctions that required further manual labor to identify the configuration problem. With the configuration data stored as XML in a single repository, the department simply has to change a configuration once, test it against the repository for conflicts, and then push it out to all appropriate infrastructure components. This eliminates configuration-related application downtime and saves an enormous amount of time as well.

Migration, not Migraines

By eliminating the manual configuration productivity drain, the IT department is well on its way to migrating all of its applications to WebSphere 5.0. However, the benefits of the solution will not end once the migration is complete. Any time a configuration needs to be touched – whether it's for deployment, retirement, or simple modification – the threat of configuration chaos looms. Additionally, automated configuration management has radically simplified the company's offshore development operations. In the past, these operations would use their own set of configurations for developing and testing applications; then, the IT department would have to reconcile those configurations with its own infrastructure – again, by hand – before the applications could be deployed.

The automated infrastructure configuration management solution enables all developers – both in-house and offshore – to work off the same, centralized repository of configurations. This allows the company to clone entire application environments to ensure total consistency across all stages of the application life-cycle and across geographic and organizational boundaries. By building sets of configuration files that specify the complete settings of all elements of the application infrastructure, and then storing these configuration sets centrally, the solution enabled the IT department to synchronize the development, testing, staging, and production environments across the U.S. and abroad.

Conclusion

In providing a single, consistent way to manage configurations across the entire application life-cycle – through development, testing, certification, and deployment – the IT department estimates that the automated infrastructure configuration management solution has slashed application migration time by nearly 60%. So, it has paid for itself in a matter of weeks. More importantly, it has substantially improved the return on investment in the WebSphere 5.0 platform, because workers can deploy more applications in a shorter period of time. They can reduce the amount of time and effort required to manage those applications. Outsourced development operations are much more tightly integrated with internal operations. And, most importantly, configuration-related downtime has virtually disappeared.

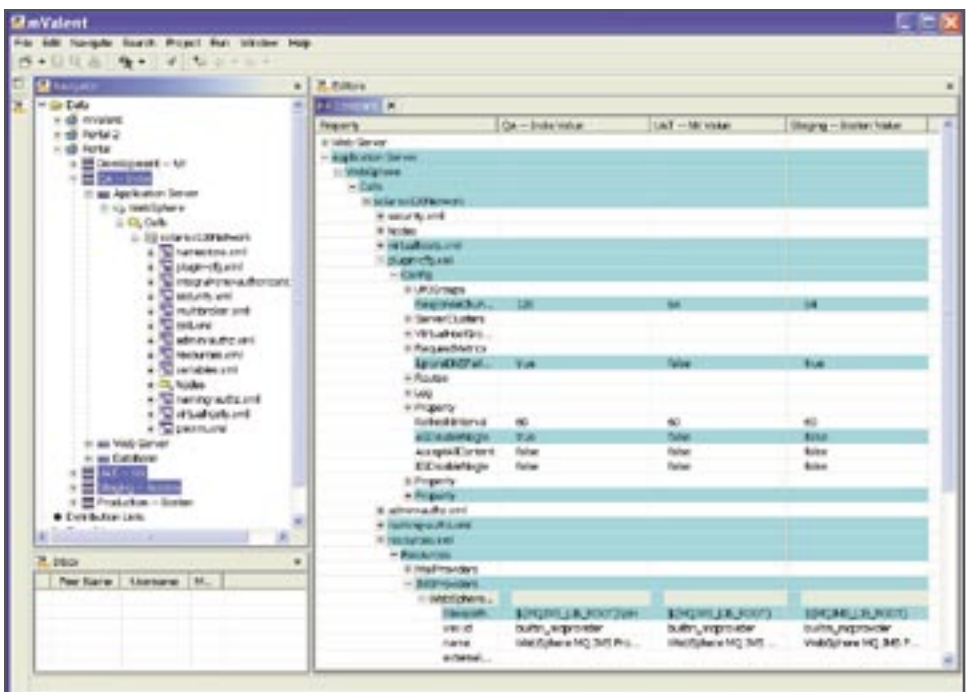

FIG 2: AUTOMATED INFRASTRUCTURE MANAGEMENT SOLUTION

Software solutions extend the functionality of the insurance industry

Updating Key Business Practices and IT Systems

BY KARLA FEUER

Karla Feuer has been a corporate communications consultant since November 1994, specializing in strategies and implementation of employee and executive communications, public relations, newsletters, Web site content, and general writing and editing. She has handled numerous assignments for IBM, including eight years with the IBM Software Group, producing newsletters, executive communications, and Web writing. swsnews@us.ibm.com

The insurance industry is facing pressures on its financial results that rival those challenging almost any other industry. On the expense side, there is increased loss – from man-made catastrophes, new environmental and liability issues, and inflation – and customer demands for better service. On the investment side, insurers' critical investment income has dropped.

At the same time, many insurers need to modernize key business processes and IT systems. They need to transform isolated applications based on products into integrated capabilities based on comprehensive customer information. Given industry conditions, projected losses, and less-than-stellar investment income, this individual product or process focus is considered neither sustainable nor affordable.

"Insurance companies need to make a cultural shift to look at their businesses from a customer view, allowing them to better up-sell and cross-sell more to generate additional revenue and cut the cost of attaining business. Additionally, customers are refocusing on their core business, specifically the underwriting discipline to better manage the business," said Cindy Maike, Global Market Segment Manager, Insurance, for IBM Software Group. "The technology exists to make this change. Insurers need to extend the functionality of their existing systems. Today, several [companies] are doing this, however, it is not a common business practice in the industry."

IBM has been a key provider of software to the insurance industry. Over a billion insurance transactions per day are handled by IBM data management products. The top 10 insurers in the world, and three-quarters of the top 50, use IBM's WebSphere products. Furthermore, IBM has the largest IT/business services organization (15,000+ dedicated) serving the insurance industry and has received over 25 insurance-specific patents since 1999.

This experience is the foundation of IBM's Industry Middleware Solutions for insurance, which are designed to address the industry's most pressing challenges. The solutions are part of IBM's effort to deliver middleware solutions based on customer preference for buying solutions designed for their industry.

Each solution contains capabilities from SWG's five brands and industry-specific middleware. The solutions are often combined with applications from independent software vendors (ISVs), as well as industry-expert services.

There are five IBM middleware solutions for insurance. They are modular, which allows companies to implement components at their own pace, and are based on open standards that support platform independence. They are based on IBM's Insurance Application Architecture (IAA), which is an evolving framework for understanding and structuring applications for the insurance industry introduced 12 years ago.

"IAA provides a way to assess an insurance company's requirements and meet them in a way that builds on their existing systems and investments," said Harsh Dave, a Solution Architect for IBM. "They can use IAA as a reference model to accelerate development, addressing one or more functional areas sequentially or in parallel, building reusable components for a services-oriented application architecture."

Flexibility also exists because of the platform independence. This enables insurers to modernize and integrate virtually every business process – within or outside of their organization – so they can respond quickly to customer needs, competitive and market challenges, and regulatory requirements. They can be on-demand businesses.

The solutions were designed to help insurers transform "large, difficult-to-maintain application environments," said Dave. "And reduce operating expenses while balancing customer satisfaction," added Maike.

The solutions include the IBM Middleware Solution for Insurance Integrated Claims Management, which provides the ability to manage the complete claims cycle. Most insurers see over 80% of premiums

flow through their claims organization with multiple interactions and many applications that have limited integration. The result is frequent process delays that erode customer satisfaction.

Using mobility, automation, and process integration, the claims management solution helps insurers improve cycle time, fraud detection, and loss ratios. It provides Web-based access to claims applications, thus extending systems to field workers; automates the scripting of the claims process; and provides skilled role-based routing of adjusters and workload balancing.

As a result, the claims management solution helps ensure the proper and timely assignment of claims adjusters and improves integration with third-party providers. Bottom-line benefits include decreased claims expense and increased customer satisfaction.

UBench, the Belgian-based provider of Web-based automotive services and solutions, needed an integrated, real-time workflow system for auto insurance claims processes. The IBM solution was used to create a Web-based platform that provided cross-enterprise collaboration and workflow business processes for claims settlements.

The result was a reduction in communication costs and in claim cycle times from weeks to days, thereby increasing customer loyalty and subscriber satisfaction. Full payback is expected in three years and the open standards–based solution provides ease of integration with legacy and third-party systems.

Progressive Insurance, the third largest provider of auto insurance in the U.S., wanted its field claims representatives to serve claimants on site. The IBM solution redefined Progressive's process for claims handling in the field, streamlining it, and cutting costs.

Insurers are also trying to cut the cost of underwriting and increase the profitability of new business. Today, most underwriters have limited integration with agencies and focus on correcting errors rather than on underwriting discipline. This causes information bottlenecks and delays processing.

IBM's Middleware Solution for Integrated Underwriting helps improve productivity and capacity by providing a collaborative workplace for underwriting activities. The solution facilitates the automation of routine application processing, including the provision of secure remote access for mobile devices.

It also allows insurers to effectively determine pending requirements – providing them with access to underwriting information to assess risk – as well as the ability to dispose of and issue policies using business-rules

based workflow tools. The underwriting solution provides for industry standard–based integration of agents and third parties and the proper and timely assignment of underwriters.

Acuity Mutual Insurance Co. of Sheboygan, Wisconsin needed an online policy rating and new business submission system integrated with its third-party rating network. Using the IBM solution, Acuity integrated its internal and external business processes. The results were a 35% increase in premiums per independent agency, $93 million in new business, 98% of new applications processed within 24 hours, and a 97% increase in its customer satisfaction rating.

Insurers also need to streamline their policy administration processes; these are often redundant and based on fragmented legacy applications that are not flexible. This makes it difficult for insurers to rapidly design, develop, and introduce new products, coverage, and rates via their choice of access devices. It also hinders efforts to respond rapidly to regulatory requirements and provide relevant time information for a variety of functions.

The IBM Middleware Solution for Insurance Policy Administration helps reduce service costs, and increase profitability and customer retention by streamlining processes for quotes, new policies, and renewals. It does so by using insurance industry standards to facilitate the integration of agents and third parties with best-fit components and by coordinating multiple types of data within a single workflow. The result is reduced process bottlenecks and secure, reliable, and robust system infrastructure and monitoring.

The California State Automobile Association (CSAA), part of the American Automobile Association, or AAA, was facing increased competition and losing market share. The association, which provides towing services, travel information, multiline insurance, and vehicle financing, needed to increase membership retention and grow revenue, but existing core systems could not support this flexibility and growth. Using IBM software, new business processes were built with an upgrade of CSAA's core insurance components.

The project is still being implemented, but after the integration of internal and external processes, it is expected to provide a single 360-degree view of the customer and real-time status of service requests.

Insurers also need to integrate and manage multiple channels while recognizing and reacting to customers at each access point. With declining premiums, insurers need to focus on profitable opportunities, improving overall customer profitability, and retaining agents. This has been difficult with redundant information, little automation, and inadequate access to customer data by the field force.

IBM's Middleware Solution for Insurance Channel Distribution Integration improves transaction efficiencies and reduces operational costs. It does so by providing agents, insurers, and re-insurers with a single-entry multiple company interface (SEMCI) and secure accessibility to information about the insured, with their method and device of choice.

Customers of Prudential wanted access to their account and benefit information over the Internet. Prudential realized early on that its technology needed to provide the strongest possible security model to minimize exposed code, protect customer data, and maintain customer trust. Using IBM software, Prudential built an integrated portal providing real-time account access, product information, and planning services.

As a result, Prudential lowered its call center costs, with a significant year-on-year reduction in call volumes. With the new system, which supports 300,000 registered users, 15% of exchanges and 5% of redemptions were processed online.

Besides responding quickly to customers with specific requests, insurers also need to develop a deeper understanding of customers' overall needs. This is critical, as they struggle to sustain profitable growth, but is currently difficult because their information is largely product-based and highly repetitive.

IBM's Middleware Solution for Insurance Customer Insight helps insurers focus on customers, reducing customer acquisition costs. It provides a single-customer view for analysis and segmentation of broad customer information, as well as campaign management and sales automation capabilities. The solution enables new business models that meet customer needs and drive new business growth.

Aetna Life Taiwan needed to track customer behavior and protect against fraudulent claims and other abuses, which cost insurers millions each year. They implemented an IBM solution that made Aetna Life "much smarter about running its business," said Miller Wang, assistant vice president and chief information officer. "For example, we're more cost-efficient in marketing the right products to the right customers, and in forecasting the prices our customers would be willing to pay."

Aetna's company-wide access to critical data improved its ability to develop products and services and to combat insurance fraud.

German insurer Provinzial had duplicative information in legacy applications and lacked integrated information about customers. This led to inconsistent marketing data inhibiting cross- and up-selling. Their IBM solution provided access across disparate warehouses and marts for a unified client view. This integrated view of individual policy information provided consistent data for management of campaigns, data warehousing, and distribution channels.

The insurance industry faces business challenges across all its processes, from marketing to claims management, yet it has a long history of bringing IT to bear on its issues. And IBM has a long history of helping it to do so.

web services conference EDGE &expo

Web Services Edge 2005 East

International Web Services Conference & Expo

Hynes Convention Center, Boston, MA

February 15-17, 2005

The Largest i-Technology Event of the Year!

Guaranteed Minimum Attendance 3,000 Delegates

Tuesday, 2/15:	Wednesday, 2/16:	Thursday, 2/17:
Conference & Expo	Conference & Expo	Conference & Expo

Join us in delivering the latest, freshest, and most proven Web services solutions... at the *Fifth Annual Web Services Edge 2005 East – International Conference & Expo* as we bring together IT professionals, developers, policy makers, industry leaders and academics to share information and exchange ideas on technology trends and best practices in secure Web services and related topics including:

- Transitioning Successfully to SOA
- Federated Web Services
- ebXML
- Orchestration
- Discovery
- The Business Case for SOA
- Interop & Standards
- Web Services Management
- Messaging Buses and SOA
- SOBAs (Service-Oriented Business Apps)
- Enterprise Service Buses
- Delivering ROI with SOA
- Java Web Services
- XML Web Services
- Security
- Professional Open Source
- Systems Integration
- Sarbanes-Oxley
- Grid Computing
- Business Process Management
- Web Services Choreography

3-Day Conference & Education Program

features:

- Daily keynotes from companies building successful and secure Web services
- Daily keynote panels from each technology track
- Over 60 sessions and seminars to choose from
- Daily training programs that will cover Web Service Security, J2EE, and ASP.NET
- FREE full-day tutorials on .NET, J2EE, MX, and WebSphere
- Opening night reception

Interested in Exhibiting, Sponsoring or Partnering?

Becoming a Web Services Edge Exhibitor, Sponsor or Partner offers you a unique opportunity to present your organization's message to a targeted audience of Web services professionals. Make your plans now to reach the most qualified software developers, engineers, system architects, analysts, consultants, group leaders, and C-level management responsible for Web services, initiatives, deployment, development and management at the regions best-known IT business address – The Hynes Convention Center in Boston.

For exhibit and sponsorship information please contact Jim Hanchrow at 201.802.3066, or e-mail at jimh@sys-con.com.

Sponsored by:

JDJ · WebServices · SD Times · wldj · .NET Journal · WebSphere Journal · Information Storage+ Security Journal · XML Journal · asp.netPRO · MX Developer's Journal · ea industry consortium

Contact for Conference Information: Jim Hanchrow, 201-802-3066, jimh@sys-con.com

Produced by SYS-CON EVENTS

www.sys-con.com/edge

Developing a comprehensive solution

Data Transformation in the WebSphere Business Integration Platform

BY JOSEPH SCHWARTZ

As vice president of marketing, Joe is responsible for Itemfield's product pricing, positioning, packaging and the formulation of the company's ongoing go-to-market strategy. Previously, he was the vice president of product management at IONA Technologies, Inc., a middleware company. Joe also founded and managed two successful technology startups, one of which went public on the NASDAQ.
josephs@itemfield.com

The most basic yet challenging problem of business integration is reconciling differences between data and messages within the enterprise. Enterprises have multiple application systems that hold data related to the data in other applications. The problem is the way in which those data are expressed – the semantic (or meaning) and syntactic (or format) incompatibilities between the disparate data models and representations used by applications. In a way, the data model for all of these different bits and pieces is inconsistent.

The crucial role of transformation in a business integration strategy is to be able to map these inconsistent data formats and data models into common objects that can be shared among the various applications. The role of data transformation is crucial, because at each and every step of a business process, transformation of the data is typically required.

The ability to transform complex data is of particular importance today because the rush to operational efficiency and regulatory compliance pressures has created the need for a new class of high-performance, mission-critical applications. These applications require accurate real-time data transformation between unstructured and semistructured document and message formats, such as nonstandard variants of EDI, legacy COBOL structures, XML, and proprietary business document types such as Microsoft Office and Adobe PDF.

Data transformation is often a major barrier to achieving business value and ROI from business integration projects, due to the diversity of formats used in the modern enterprise. Until now, data transformation has been accomplished with tools and technology that:

- Fail to address diverse data types and often require costly custom development
- Exacerbate technology proliferation issues
- Drive up overall integration costs, complexity, and project timelines

CLASSES OF TRANSFORMATION PROBLEMS

From a business viewpoint, the most pressing problems associated with data transformation are:

- The need to channel a large and diverse amount of information to end users and systems in support of customer self-service and real-time enterprise applications/initiatives such as B2B and other "extended enterprise" connections.
- Regulatory compliance requirements such as Sarbanes-Oxley, HIPAA in healthcare, the U.S. Patriot Act, and Basel II accords in financial services. Compliance imposes control, reporting, and data integration challenges.

And, of course, another significant problem in these cases is cost. On average, application integration implementation costs range from three to five times the cost of required software licenses. Data transformation accounts for a disproportionate share of that implementation cost, which grows even higher as the data integration requirements increase in complexity.

And the problems will not soon go away. Unstructured data accounts for as much as 75% of the information stored in enterprises today. Semistructured data, which includes a variety of B2B and legacy data formats, is prevalent in many business integration scenarios and business application initiatives, such as B2B transaction integration, collaborative document-oriented workflow, and customer-self service.

TECHNICAL REQUIREMENTS FOR TRANSFORMATION

Recent standards such as XML, Extenible Style Sheet Language (XSL), XSL Transformation (XSLT), and the J2EE Connector Architecture (J2CA or J2EE-CA), bring some order to the world of data transformation. However, they don't address the full set of data transformation requirements for the business integration challenges that enterprises need to solve today.

In addition to robust support for the conversion of multiple data formats to and from XML, which is rapidly becoming the lingua franca for business integration, a comprehensive solution for transformation must support the integration of complex data and should include:

- Prebuilt functions for one-to-many and many-to-many transformations
- Support for transformation of repeating groups
- An intuitive graphical development environment for designing and maintaining transformation maps (because even standard data formats may change all too quickly)
- Support for custom transformation functions
- Support for processing sets of data (e.g., dealing with multiple input records and associated filtering, summarizing, and sorting)
- Support for unstructured and semistructured documents
- Automatic discovery/import of document metadata – the ability to create definitions from a format specification or an example document

DEVELOPMENT ACTIVITIES IN TRANSFORMATION

From the technical perspective, the transformation process involves the following activities:

- ***Sourcing:*** finding and understanding the structure of the candidate data specification or data instance for transformation
- ***Definition:*** specifying the information to be extracted from the source, which is difficult to automate for complex data without some form of data visualization
- ***Mapping:*** the specification of the semantic relationship between source and target formats, which is difficult to automate for complex data that typically doesn't have the hierarchical structure required by most mapping tools
- ***Translation:*** the implementation of syntactic changes applied to the data
- ***Reconciliation:*** the validation of the transformation; deals with any inconsistencies in the data

Today, these activities are accomplished using inefficient tools and processes that rely heavily on custom programming. What's needed is a rational and comprehensive approach to data transformation that addresses the greatest inefficiency in the process: the programming – and maintenance – of intensive transformation activities.

Data Transformation in the WebSphere Business Integration Platform

The IBM Business Integration Reference Architecture (BIRA) encompasses the key areas of integration capability required for comprehensive, enterprise-wide integration strategies and solutions. In the context of the BIRA, scenarios are presented in areas where the need for a robust data transformation solution is most apparent. These scenarios also mention the specific WebSphere Business Integration components that carry out complex data transformations, and those components are then described.

BUSINESS INTEGRATION REFERENCE ARCHITECTURE SCENARIOS

Figure 1 shows the IBM Business Integration Reference Architecture, including both the generic services and the IBM products that deliver the required generic capabilities.

When automating processes across the value chain, data is inevitably encountered that isn't self-describing and that is not defined with schemas, as are XML documents. To integrate across systems, such unstructured data must be accommodated. Below are several scenarios that map to various BIRA capabilities and the data transformation requirements they entail:

- ***Partner Services:*** These are provided by WebSphere Business Integration Connect for the exchange of B2B transactions. These transactions may involve standards-based data, and custom implementations of those standards, or the exchange of binary documents.
- ***Application and Data Access:*** For application and data access services, WebSphere Business

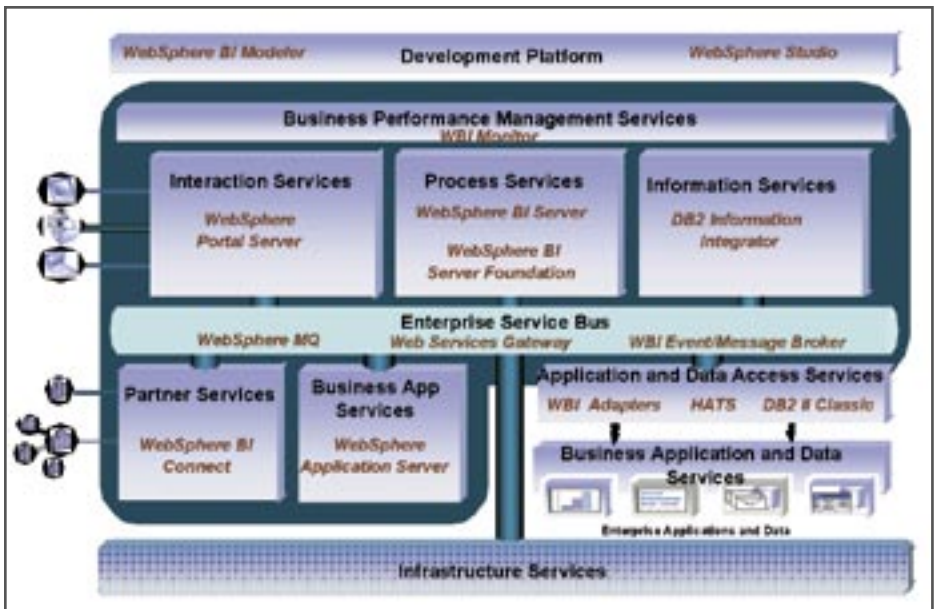

FIG 1: IBM BUSINESS INTEGRATION REFERENCE ARCHITECTURE

Integration Adapters exchange transactions in their native application formats, enabling applications to participate in a service-oriented architecture (SOA). In many cases, transforming complex application formats will be required.

- ***Enterprise Service Bus (ESB):*** Within the Enterprise Service Bus, which transports and mediates transactions across distributed systems in an SOA, ESB clients may exchange data in complex data formats. In this case, WebSphere MQ, Web Services Gateway, and the WebSphere Business Integration Message Broker are used.
- ***Process Services:*** For process services, where solutions implement multistep business processes that span people and systems, complex data must be also be accommodated. In this case, the WebSphere Business Integration Server and WebSphere Business Integration Server Foundation are used.

DATA TRANSFORMATION TOOLS IN WEBSPHERE BUSINESS INTEGRATION COMPONENTS

This section lists the WebSphere Business Integration Platform components and the tools they provide for data transformation:

- WebSphere Business Integration Connect enables a business to extend integration beyond the enterprise – exchanging information with trading partners and managing trading communities. WBI Connect can use schema validations, XSLT transformations, and custom exits to call external transformation services on the transmitted documents.
- WebSphere Business Integration Server provides process integration, workforce management, and enterprise application connectivity. It has the following sub-components:
  - WebSphere InterChange Server automates and synchronizes business activities executed across multiple applications as business processes. Its GUI-based Map Designer defines and generates code for transformation maps between application-specific and normalized data models. It specifies the transformation steps for each destination attribute to be transformed. Mapping is supported between data values and data entity relationships, and via access to third-party mapping products and databases.
  - WebSphere MQ Workflow enables the integration of all participants in the business process, including those external to the organization. Data mapping tools are provided via the WebSphere MQ Workflow Buildtime graphical environment.
  - WebSphere Business Integration Message Broker transforms and enriches in-flight information to provide a level of intermediation between applications that use different message structures and formats. It provides rich support for configuring message flows that transform messages from one format to another with the Message Flow Mapping editor in the Message Brokers Toolkit for WebSphere Studio.

For messages that are pre-defined – their content and structure are both known and predictable – WebSphere Business Integration Message Broker can use the facilities provided by the MRM. For messages that aren't predefined, a common approach is to include a Compute node in the message flow to create the output message with the required content. This task can vary greatly in complexity and effort, depending on the complexity of the format. For example, for each compute node in a given flow or application, an ESQL module must be coded using ESQL statements and functions to tailor the behavior of the node.

From a runtime perspective, WebSphere Business Integration Message Broker includes prebuilt parsers. Parsers are programs that interpret the bit streams of incoming messages and create internal representations of the message.

  - WebSphere Business Integration Server Foundation is a standards-based platform optimized for building and deploying composite applications by creating reusable services out of Web services, Java assets, back-end systems, and packaged applications. It includes WebSphere Process Choreographer, which is a business-process engine that allows for the efficient execution of business processes. With it, business process technology can be combined with any other service offered by the open J2EE architecture.
  - WebSphere Business Integration Server Foundation includes WebSphere Studio Application Developer Integration Edition, which provides a drag-and-drop integration development environment optimized for building composite applications. Studio provides a number of XML tools for building and validating DTDs, XML schemas, and XML files; for generating JavaBeans from a DTD or XML schema; and for defining transformation mappings between XML documents by generating XSLT scripts. In addition, Studio tools support the creation of an HTML document by applying an XSL style sheet to an XML document, the definition of mappings between relational tables and DTD files, or between SQL statements and DTD files.

• WebSphere Business Integration Adapters enable an enterprise to create integrated processes that exchange information between ERP, HR, CRM, and supply chain systems. There are several different types of adapters as well as a toolkit for developing custom adapters:

- ***Applications Adapters:*** extract data and transaction information from cross-industry and industry-specific packaged applications, and connect them to a central hub
- ***Mainframe Adapters:*** provide access to application data in OS/390 systems and provide connectivity approaches to AS/400 systems
- ***Technology Adapters:*** provide the connectivity for accessing data, technologies, and protocols that enhance integration infrastructure

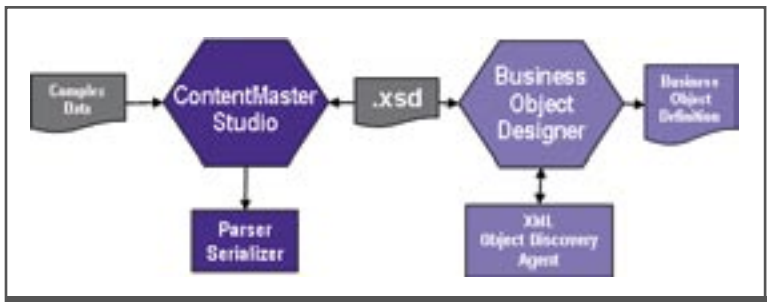

FIG 2: CONTENTMASTER AT DESIGN TIME

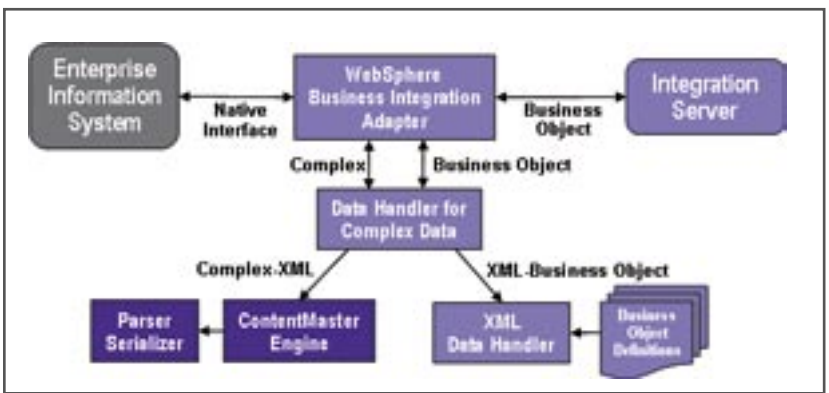

FIG 3: CONTENTMASTER RUNTIME CONFIGURATION

WebSphere Business Integration Adapters use data handlers to perform data transformations and to manage interactions with both WebSphere Business Integration Platform components and applications. IBM provides several standard data handlers (e.g., Fixed-Width, Delimited, and Name-Value data handlers) and special data handlers (XML and EDI data handlers).

The WebSphere Business Integration Adapter Framework enables adapters to be used in many different Business Integration solutions, including those provided by WebSphere InterChange Server, WebSphere Business Integration Message Broker, or WebSphere Business Integration Server Foundation.

ADVANCED DATA TRANSFORMATION

For the integration of complex data, IBM provides the WebSphere Business Integration Data Handler for Complex Data. It allows for the bidirectional conversion of text and binary formats to and from WebSphere Business Integration business objects. By using this data handler, a user can integrate with standard and proprietary formats such as Microsoft Word, Microsoft Excel, Adobe PDF, COBOL Copybooks, and HL7.

This data handler is commonly used with a WebSphere Business Integration Technology Adapter, e.g., JText, HTTP, MQ, or e-mail. It can also be used with a custom-built adapter. The IBM data handler leverages the Itemfield ContentMaster product, which provides a visual development and test environment for constructing custom parsers, without requiring custom programming.

Custom Java data handlers can be created that manipulate the Data Handler API by using the Adapter Development Kit and writing custom Java code. However, when data integration complexity warrants coding of a custom data handler, the IBM/Itemfield solution offers an excellent alternative.

The WebSphere Business Integration Data Handler for Complex Data implementation leverages XML to exchange data between the data handlers and Itemfield ContentMaster.

DESIGN TIME TOOLING

Figure 2 shows a design time implementation in which ContentMaster Studio is used to generate the parsers and serializers needed to deal with the complex data stream. An XML schema is used to describe the XML exchange with the data handler.

The XML schema is then utilized by the WebSphere Business Integration Adapter tooling to generate business object definitions describing the data exchanged between the adapter and the integration server.

RUNTIME INFRASTRUCTURE

Figure 3 illustrates runtime configuration. Here, the adapter manages the exchange of business objects with the integration server, as well as the exchange of data streams with the application, using its native interface. Using the data handler, it translates between data streams and business objects. It does so by leveraging the embedded ContentMaster Engine with its serializers and parsers, as well as the XML data handler.

Because the Data Handler for Complex Data can be plugged into a technology adapter, ContentMaster can be used in virtually all WebSphere Business Integration components. The adapter framework is accessible throughout the entire WebSphere

–continued on page 48

Adopting the right techniques for the right J2EE application

A Practical Guide to J2EE Development

BY THINA NATARAJAN

Thina Natarajan is a senior solutions architect in IBM's Tivoli Software Group. He optimizes IT architecture, development, and deployment phases for global clients. Recently, he has been focusing on J2EE performance tuning best practices.
thina@us.ibm.com

Evolving J2EE specifications continue to provide developers and architects with added performance muscle and flexibility for building and optimizing enterprise applications. J2EE advances are both a boon and a challenge as they can provide additional capabilities while also requiring refined development skills. J2EE applications incorporate complex distributed logic and, hence, demand careful, intelligent, and innovative design and implementation techniques. I'll outline various programming design techniques that empower J2EE developers to create more robust enterprise applications.

A J2EE application is comprised of multiple artifacts, known as Java components, which act like musicians in an orchestra. Each J2EE artifact is required to perform its role, while working in concert with the other components to create the complete J2EE symphony. Developers and architects serve as composers, with the task of ensuring that each component "instrument" plays its assigned notes – without causing performance discord by consuming too many memory resources or failing to deliver its artifact "notes" on queue. The large number of instruments in a J2EE symphony makes it difficult for developers and architects to identify the source of out-of-tune components. J2EE architects and developers, therefore, play a vital role in aligning individual artifacts into powerful and scalable J2EE applications.

The adoption of proven design and implementation techniques can help ensure that J2EE applications are flexible, scalable, and can be used as templates for new projects. The following sections outline best practices that every developer or architect "composer" should use to support successful J2EE implementation.

Choose the Best-Fit Application Architecture

There's no single application architecture that addresses all organizational goals. As such, organizations must use business objectives and end-user requirements as foundation footings for their application infrastructures. An architecture should enable the application to scale gracefully and have the performance to virtually eliminate Mean Time Between Failures (MTBF) and Mean Time To Recovery (MTTR). The architecture should also take advantage of object-oriented (OO) design patterns and principles, which support object scope extensibility and promote reusability, portability, and parsimonious use of computing resources.

Perhaps the biggest question posed by J2EE implementations is the use (or misuse) of Enterprise JavaBeans (EJBs). EJBs constitute a major and complex component of the J2EE specifications and remain a focal point for application server vendors. Many distributed Web applications maintain the state and session information of their clients and connections. It's essential that organizations select a data access methodology – such as direct Java Database Connectivity, Data Access Objects (DAO), or Java Data Objects (JDO) – that addresses both application user requirements and overarching business goals.

In addition to taking full control of the session life cycle, application architectures should include proven procedures for storing, restoring, recovering, and releasing state information. These decisions should be independent of application deployment patterns such as cloning, clustering, and split Java Virtual Machines (JVMs). For example, while state information can be stored either at the Web tiers or at the EJB tier, the application architecture should target the Web container to store the state information.

Enterprise design patterns should be the defining yardsticks for application architectures. Modeling and

implementing business objects using design patterns creates a productive development environment that is built on reuse and harvests proven knowledge and assets. The J2EE specification necessitates intricate handling of checked and unchecked exceptions. Application architectures, therefore, should enable developers to identify and isolate exceptions, present a meaningful message to the presentation layer, and provide a straightforward recovery path to the business tier.

Prevent Objects From Flooding the JVM

The adage "a house is only as strong as its foundation," can also be used to underscore the role of JVMs as the bedrock of a J2EE application. JVMs are arguably the foundations for applications; J2EE components are built and reside in JVMs. The majority of J2EE application performance problems are caused by JVM heap allocations failures, frequent garbage collections, and out-of-memory situations.

Developers can sidestep the majority of J2EE application performance challenges by using an application architecture that judiciously creates new objects, emphasizes object reuse, and extends application relevance through additional responsibility (behavior) and a wider scope in the domain. This approach significantly enhances JVM performance by maintaining smaller heap sizes and fewer object allocations, which in turn minimize garbage collections and eliminate out-of-memory situations that cause application performance problems.

The developer community is largely unaware of Java's robust heap management and garbage control capabilities. The java.lang.ref package, for example, classifies object references into active, passive, convenient, and dead references. The package also enables developers to manipulate objects through event queues and listeners. The java.lang.ref enables business objects to be wrapped in different reference types, which gives the JVM greater control over object life spans. Java.lang.ref requires a deep understanding of the garbage collection process. Once Java.lang.ref is understood and used, however, developers can create code that optimizes runtime memory utilization.

Java Development Skills That Complement J2EE Design Skills

J2EE architectural excellence without programmatic prudence is an exercise in futility. Many projects do not realize their potential because elegantly crafted Java code is compromised by poor application design. Similarly, many applications have failed to deliver desired performance because developers strayed from the rigid coding disciplines required to address the complexities of J2EE application program interfaces (APIs). A typical J2EE application contains distributed logic, which requires a balanced approach to address myriad variables, including algorithms, concurrency, isolation, transaction management, and inter-tier communication. Code development for APIs, in turn, must address resource conservation, include sound exception-handling features, and allow graceful application recovery.

Some major Java coding challenges that developers face, include implementations that fire multiple-threads, hard coding of transactions, programmatic control of security roles and persistence, tight coupling objects between tiers, and caching huge

WebSphere Journal Coming Next Month...

Can Your Car Talk?

And do you know what it's saying?

AN INTERVIEW WITH BRUCE RADLOFF, CHIEF TECHNOLOGY OFFICER OF ONSTAR, AND TONY LENT, VICE PRESIDENT OF ONSTAR'S WIRELESS STRATEGIC BUSINESS UNIT

The World of Telematics

Connecting vehicles to back-end systems

AN INTERVIEW WITH JIM RUTHVEN, PROGRAM DIRECTOR FOR IBM TELEMATIC SOLUTIONS

JavaServer Faces (JSF) vs. Struts

A brief comparison

BY ROLAND BARCIA

Streamlining Core Business Processes

IT solutions for the financial market industry

BY DAVID LAPP

amounts of information for each client thread. Many developers also add needless strain to their Java environments because they create redundant features in J2EE applications that are already available in J2EE API and application server environments. Many developers, for example, include caching features in J2EE applications that are already available in the underlying J2EE application servers. A thorough understanding of J2EE specifications and APIs, as well as a sound awareness of the extended features provided by the application servers, will yield higher productivity with quality and optimized resources. Good J2EE architecture and developer craftsmen will continually research the best techniques and methodologies for future application initiatives.

Selectively Implement OO Design Principles

Contrary to common belief, prolific use of OO techniques does not necessarily make J2EE applications more robust. Deep inheritance hierarchies, in which each class has a large number of ancestors, can often be counterproductive because they expand the JVM object stack and require increased overhead in object creations. J2EE containers rely heavily on object serialization for session persistence management; serialization of such objects with deep inheritance is a strain on the application server.

Object hierarchy adds overhead unless carefully handled through techniques such as declaring class variables as transient and implementing externizable interfaces. Flattening an object hierarchy requires a balance of performance and maintainability during development. Nevertheless, a flatter object hierarchy minimizes object allocation and collection, which supports optimal performance. Programming to interfaces – not classes – and implementing façade patterns enables loose object coupling, adds more flexibility, allows for coarse-grained information exchange, and decreases network latency. Developers reduce design complexity and optimize run-time garbage collection by declaring class constructs, such as instance variables and method arguments, to the least specific type (i.e., the highest ancestors in the hierarchy).

Gang of Four (GoF) patterns, particularly strategy and state design patterns, favor object compositions over class (concrete) inheritance. This is a particularly suitable alternative to class inheritance because it provides developers with the added flexibility to alter an object's behavior during run time through delegation. An additional way of achieving flexibility is programming to interfaces rather than to classes. Interface-based architectures add strength and scalability to J2EE applications because of their simplicity and loose coupling between objects.

An application architecture that minimizes object allocations, increases the life and scope of objects, and supports the effective reuse of objects through reinitialization and pooling will perform well. Multi-threading, for example, can be avoided by using asynchronously concurrent Message-Driven Beans (MDBs), which demand less from the JVM, avoid complexity, and reduce object count. Finally, synchronization of code blocks should be justified and avoided as much as possible.

J2EE software development is a complex, multifaceted process, much like the business cases it solves. The four main points of this article are intended to serve as seeds for discussion and debate among J2EE development teams when determining best practices for architectural and development strategies. It's the responsibility of architects and developers to harness the potential of J2EE for their organizational objectives.

References

- Stearns, Beth, et al. (2002). *Designing Enterprise Applications with the J2EE Platform, Second Edition*. Addison-Wesley.
- Johnson, R. (2003). *Expert One-on-One J2EE Design and Development*. John Wiley & Sons, Inc.
- Crupi, John, et al. (2001). *Core J2EE Patterns: Best Practices and Design Practices*. Prentice Hall.

–continued from page 45

Business Integration stack.

Conclusion

The WebSphere Business Integration Platform provides many options for transforming data, including off-the-shelf adapters and data handlers. However, dealing with complex data is the heart of today's data transformation issue. These off-the-shelf components may not provide all of the necessary support for creating custom transformation logic that the integration platform doesn't provide out of the box, or for handling input formats that require too much maintenance.

For the integration of complex data, IBM has provided a compelling solution that provides an attractive alternative to developing custom Java data handlers within adapters. Complex data transformations can be implemented rapidly, without custom programming, making them very flexible in the face of change. IBM's tight integration of Itemfield's next-generation data transformation solution extends the scope and power of the WebSphere Business Integration platform by automating the integration of complex data – the most challenging part of many business integration efforts.

A LIMITED TIME SAVINGS OFFER FROM *SYS-CON MEDIA*

SUBSCRIBE TODAY TO MULTIPLE MAGAZINES

AND SAVE UP TO $340 AND RECEIVE UP TO 3 FREE CDs!*

RECEIVE YOUR DIGITAL EDITION ACCESS CODE INSTANTLY WITH YOUR PAID SUBSCRIPTIONS

CALL TODAY! 888-303-5282

Pick a 3-Pack, a 6-Pack or a 9-Pack

☐ 3-Pack	☐ 1YR	☐ 2YR	☐ U.S.	☐ Can/Mex	☐ Intl.
☐ 6-Pack	☐ 1YR	☐ 2YR	☐ U.S.	☐ Can/Mex	☐ Intl.
☐ 9-Pack	☐ 1YR	☐ 2YR	☐ U.S.	☐ Can/Mex	☐ Intl.

TO ORDER

•Choose the Multi-Pack you want to order by checking next to it below. •Check the number of years you want to order. •Indicate your location by checking either U.S., Canada/Mexico or International. •Then choose which magazines you want to include with your Multi-Pack order.

☐ **LinuxWorld Magazine**

U.S. - Two Years (24) Cover: $143	You Pay: $79.99 /	Save: $63 + FREE $198 CD
U.S. - One Year (12) Cover: $72	You Pay: $39.99 /	Save: $32
Can/Mex - Two Years (24) $168	You Pay: $119.99 /	Save: $48 + FREE $198 CD
Can/Mex - One Year (12) $84	You Pay: $79.99 /	Save: $4
Int'l - Two Years (24) $216	You Pay: $176 /	Save: $40 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $99.99 /	Save: $8

☐ **JDJ**

U.S. - Two Years (24) Cover: $144	You Pay: $99.99 /	Save: $45 + FREE $198 CD
U.S. - One Year (12) Cover: $72	You Pay: $69.99 /	Save: $12
Can/Mex - Two Years (24) $168	You Pay: $119.99 /	Save: $48 + FREE $198 CD
Can/Mex - One Year (12) $120	You Pay: $89.99 /	Save: $40
Int'l - Two Years (24) $216	You Pay: $176 /	Save: $40 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $99.99 /	Save: $8

☐ **Web Services Journal**

U.S.- Two Years (24) Cover: $168	You Pay: $99.99 /	Save: $68 + FREE $198 CD
U.S. - One Year (12) Cover: $84	You Pay: $69.99 /	Save: $14
Can/Mex - Two Years (24) $192	You Pay: $129 /	Save: $63 + FREE $198 CD
Can/Mex - One Year (12) $96	You Pay: $89.99 /	Save: $6
Int'l - Two Years (24) $216	You Pay: $170 /	Save: $46 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $99.99 /	Save: $8

☐ **.NET Developer's Journal**

U.S. - Two Years (24) Cover: $168	You Pay: $99.99 /	Save: $68 + FREE $198 CD
U.S. - One Year (12) Cover: $84	You Pay: $69.99 /	Save: $14
Can/Mex - Two Years (24) $192	You Pay: $129 /	Save: $63 + FREE $198 CD
Can/Mex - One Year (12) $96	You Pay: $89.99 /	Save: $6
Int'l - Two Years (24) $216	You Pay: $170 /	Save: $46 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $99.99 /	Save: $8

☐ **Information Storage + Security Journal**

U.S. - Two Years (24) Cover: $143	You Pay: $49.99 /	Save: $93 + FREE $198 CD
U.S. - One Year (12) Cover: $72	You Pay: $39.99 /	Save: $39
Can/Mex - Two Years (24) $168	You Pay: $79.99 /	Save: $88 + FREE $198 CD
Can/Mex - One Year (12) $84	You Pay: $49.99 /	Save: $34
Int'l - Two Years (24) $216	You Pay: $89.99 /	Save: $126 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $59.99 /	Save: $48

☐ **Wireless Business & Technology**

U.S. - Two Years (12) Cover: $120	You Pay: $49.00 /	Save: $71 + FREE $198 CD
U.S. - One Year (6) Cover: $60	You Pay: $29.99 /	Save: $30
Can/Mex - Two Years (12) $120	You Pay: $69.99 /	Save: $51 + FREE $198 CD
Can/Mex - One Year (6) $60	You Pay: $49.99 /	Save: $10
Int'l - Two Years (12) $120	You Pay: $99.99 /	Save: $20 + FREE $198 CD
Int'l - One Year (6) $72	You Pay: $69.99 /	Save: $2

☐ **MX Developer's Journal**

U.S. - Two Years (24) Cover: $143	You Pay: $49.99 /	Save: $93 + FREE $198 CD
U.S. - One Year (12) Cover: $72	You Pay: $39.99 /	Save: $32
Can/Mex - Two Years (24) $168	You Pay: $79.99 /	Save: $88 + FREE $198 CD
Can/Mex - One Year (12) $84	You Pay: $49.99 /	Save: $34
Int'l - Two Years (24) $216	You Pay: $89.99 /	Save: $126 + FREE $198 CD
Int'l - One Year (12) $108	You Pay: $59.99 /	Save: $48

☐ **ColdFusion Developer's Journal**

U.S. - Two Years (24) Cover: $216	You Pay: $129 /	Save: $87 + FREE $198 CD
U.S. - One Year (12) Cover: $108	You Pay: $89.99 /	Save: $18
Can/Mex - Two Years (24) $240	You Pay: $159.99 /	Save: $80 + FREE $198 CD
Can/Mex - One Year (12) $120	You Pay: $99.99 /	Save: $20
Int'l - Two Years (24) $264	You Pay: $189 /	Save: $75 + FREE $198 CD
Int'l - One Year (12) $132	You Pay: $129.99 /	Save: $2

☐ **WebSphere Journal**

U.S. - Two Years (24) Cover: $216	You Pay: $129.00 /	Save: $87 + FREE $198 CD
U.S. - One Year (12) Cover: $108	You Pay: $89.99 /	Save: $18
Can/Mex - Two Years (24) $240	You Pay: $159.99 /	Save: $80 + FREE $198 CD
Can/Mex - One Year (12) $120	You Pay: $99.99 /	Save: $20
Int'l - Two Years (24) $264	You Pay: $189.00 /	Save: $75
Int'l - One Year (12) $132	You Pay: $129.99 /	Save: $2

☐ **PowerBuilder Developer's Journal**

U.S. - Two Years (24) Cover: $360	You Pay: $169.99 /	Save: $190 + FREE $198 CD
U.S. - One Year (12) Cover: $180	You Pay: $149 /	Save: $31
Can/Mex - Two Years (24) $360	You Pay: $179.99 /	Save: $180 + FREE $198 CD
Can/Mex - One Year (12) $180	You Pay: $169 /	Save: $11
Int'l - Two Years (24) $360	You Pay: $189.99 /	Save: $170 + FREE $198 CD
Int'l - One Year (12) $180	You Pay: $179 /	Save: $1

☐ **WLDJ**

U.S. - Four Years (24) Cover: $240	You Pay: $99.99 /	Save: $140 + FREE $198 CD
U.S. - Two Year (12) Cover: $120	You Pay: $49.99 /	Save: $70
Can/Mex - Four Years (24) $240	You Pay: $99.99 /	Save: $140 + FREE $198 CD
Can/Mex - Two Year (12) $120	You Pay: $69.99 /	Save: $50
Int'l - Four Years (24) $240	You Pay: $120 /	Save: $120 + FREE $198 CD
Int'l - Two Year (12) $120	You Pay: $79.99 /	Save: $40

*WHILE SUPPLIES LAST. OFFER SUBJECT TO CHANGE WITHOUT NOTICE

3-Pack

Pick any 3 of our magazines and save up to $210.00

Pay only $99 for a 1 year subscription plus a FREE CD

- 2 Year – $179.00
- Canada/Mexico – $189.00
- International – $199.00

6-Pack

Pick any 6 of our magazines and save up to $340.00

Pay only $199 for a 1 year subscription plus 2 FREE CDs

- 2 Year – $379.00
- Canada/Mexico – $399.00
- International – $449.00

9-Pack

Pick 9 of our magazines and save up to $270.00

Pay only $399 for a 1 year subscription plus 3 FREE CDs

- 2 Year – $699.00
- Canada/Mexico – $749.00
- International – $849.00

Subscribe Online Today www.sys-con.com/2001/sub.cfm

SYS-CON MEDIA

Intersperse Announces Intersperse Manager 3.0

(Pasadena, CA) – Intersperse, Inc., a provider of management solutions for service-oriented enterprise applications, has released Intersperse Manager 3.0. Targeted at the needs of next-generation J2EE and SOA-based applications, it gives large organizations performance advances and new levels of scalability. Intersperse Manager 3.0 allows users to continually and effectively monitor the health of their enterprise IT environments.

This software offers users a consolidated view, providing whole-system visibility and real-time control over an enterprise's distributed or service-oriented application environments. By monitoring up and down the application stack and across integrated business applications, Intersperse Manager 3.0 delivers visibility, context, and manageability of system availability and performance. With it, IT departments can control tens of thousands of objects and process hundreds of events per second with one server.

www.intersperse.com

IBM Releases Software to Help Companies Automate Operations

(Somers, NY) – IBM has announced software that will enable companies to establish a more detailed roadmap of their business operations. With a better view of how business processes are being executed across disparate networks, companies will have the ability to detect slowdowns, look for efficiencies, and respond faster to customer demands and changing market conditions.

The software, WebSphere Business Integration Modeler version 5.0, uses industry standards together with integration software and development tools to help companies model the flow of business processes across their IT systems. It addresses their desires to reduce time, cost, and complexity as they engineer the increasingly automated processes.

WebSphere Business Integration Modeler allows companies to visualize, design, and chart the various operations of their organizations. It provides a means for evaluating performance and modifying processes in which slowdowns or flaws are detected.

The software is based on BPEL (Business Process Execution Language), an XML-based open standard for defining how to combine Web services to implement business processes. It provides support for WebSphere Business Integration Server Foundation, WebSphere MQ message queuing software, and Rational Rose XDE development tools. The open architecture allows customers to work with existing content that has been based on standards like XML or product content like Microsoft Visio, and extend it using WebSphere Business Integration Modeler's simulation and modeling capabilities.

www.ibm.com/ondemand

Jordanian STS Group Raises Game for IBM Portal Offerings

Specialized Technical Services Group (STS), a Middle Eastern hardware/software services and solutions provider, has been named as an IBM Premier Software Partner for the region.

Becoming an IBM Premier Business Partner involves demonstrating high levels of accomplishment in competency, commitment to IBM, customer satisfaction, and consistent revenue growth. STS BusinessONE, the e-business software division of STS, has deployed WebSphere in its secure e-payment solution. The solution, PayONE, enables debit, credit card, and bank account holders to make online payments for e-commerce and e-government transactions safely in real time. PayONE has also been implemented in the national payment gateways of Qatar Central Bank and Visa Jordan Card Services.

www.sts.com.jo

IBM Unveils Software to Help Businesses Better Manage Product Information

(Somers, NY) – IBM has developed software to help businesses better manage and gain control over product information and to enable increasingly popular initiatives such as radio frequency identification (RFID), global data synchronization, electronic commerce, and portals.

WebSphere Product Center version 5 is middleware designed to help customers develop a central, accurate, and comprehensive repository of product information that is synchronized with internal systems and shared across supply chains.

WebSphere Product Center enables customers to manage and link comprehensive product, trading partner, location, organization, pricing, and promotion information. It improves how information is handled throughout an organization by reducing errors, redundant data, and manual processes.

This software contains technology to help improve productivity and performance in global implementations that involve users, trading partners, and products.

WebSphere Product Center version 5 includes WebSphere Application Server, DB2 Information Integrator to store and access product information from any repository, WebSphere MQ, and WebSphere Business Integration Adapter for WebSphere MQ to move information between disparate systems. It has been designed to reduce both complexity and total cost of ownership for customers investing in product information management.

www.software.ibm.com

Jumping 2 Enormous Eggplants.

A new way to look at J2EE.

Cyanea/One gives you unprecedented visibility into mission-critical application performance. With the transaction-based performance management capabilities of Cyanea/One, you can drill down to the method level to determine and resolve performance issues. Cyanea/One monitors complex, multi-platform, service-oriented environments 24 x 7 from a single web-based console, all without having to touch the application code. If you are developing enterprise applications, you know the hurdles you sometimes have to overcome; with Cyanea/One, your applications will reach new heights.

For more information, visit www.cyanea.com/j2ee or call 1.877.CYANEA8.

Cyanea®