SYS-CON MEDIA

PERFORMANCE • GLOBALIZATION • INTEGRATION

WebSphere® JOURNAL

www.SYS-CON.com/WebSphere

The World's Leading Independent WebSphere Resource

MARCH VOLUME 3 ISSUE 3

TODAY'S MISSION: INTEGRATION — FULL SPEED AHEAD

KEY PRACTICE SKILLS:

WEBSPHERE APPLICATION SERVER

WEBSPHERE PORTAL SERVER

WEBSPHERE MQ

MQSERIES WORKFLOW

WEBSPHERE MQ INTEGRATOR

WEBSPHERE BUSINESS INTEGRATOR & CROSSWORLDS

WEBSPHERE COMMERCE SUITE

WEBSPHERE STUDIO APPLICATION DEVELOPER

VISUALAGE FOR JAVA

WEB SERVICES

DB2 & INFORMIX

PLANNING YOUR WEBSPHERE, WEB SERVICES OR PORTAL STRATEGY?

UPGRADING YOUR IT INFRASTRUCTURE?

NEED A JUMP START ON WEBSPHERE 5.0 AND WEBSPHERE PORTAL 4.1?

Prolifics' Jump Start Solutions will...

Educate. As 1 of 5 WebSphere service providers retained by IBM, Prolifics' Consultants are top-ranked by IBM in delivering WebSphere Solutions. We give your team an edge by transferring our extensive knowledge and experience.

Enable. We prepare your e-infrastructure for ease of application integration, migration and deployment. We recognize risks and delays that could hinder your project timeline and budget.

Empower. Providing your team with best practices, documentation and guidelines, we identify the best strategy leading to your must-meet business goals.

PROLIFICS' JUMP START SERVICES MENU:

- **Jump Start:** Extend your enterprise with WebSphere Portal 4.1
- **Jump Start:** Build, publish, manage, optimize with WebSphere 5.0
- **Jump Start:** Get started with Site Analyzer
- **Jump Start:** Integrate WebSphere MQ into a J2EE environment
- **Jump Start:** Web services – build to integrate

Call 800.675.5419 or 212.267.7722

http://www.prolifics.com

Prolifics

IBM Business Partner

Solving Business Problems with Technology Solutions since 1978
Expert Delivery of WebSphere, Integration and Business Transformation Solutions

For a FREE white paper on how to Integrate Enterprise Resources Using Web Services Technologies
Go to: http://www.prolifics.com/webservices

e-business

PRIORITY CODE: DJ0702JS

diagnosys

Learn how to improve your Web applications at H&W's free Web seminar.
Visit www.hwcs.com/wsj1.asp

You might not even know you have a leak, but that little drip can turn into a big deal if it's left unattended. Performance problems with your Web applications work the same way.

Hung threads, memory leaks, and faulty network connections can cause your application to perform poorly or even fail when it matters most – in production. Customers click away and profits plummet. But how do you find the problems if you don't even know they are there?

DiagnoSys is a complete software solution that allows you to analyze the performance of your Web applications through all phases of the life cycle. With true statistical correlation, in-depth analysis, and root cause identification, it helps your teams work together to find problems while they are still little.

The end result? With DiagnoSys, you build vital applications that optimize uptime and keep the bottom line healthy. No problem.

H&W

www.hwcs.com | 1.800.338.6692

Bridging the Enterprise Computing Gap Since 1979

IBM Offers New Role-Based Tools

BY JACK MARTIN

Well, it's another month and we're off to a good start with IBM making it easier for people to find new uses for WebSphere.

IBM has taken the teams-based approach to the next step. They just released a roles-based group of tools – six power packs and eight tracks in total – based on what you do, with two for developers focusing on either J2EE or .NET. System analysts and architects who need to model are also included. Testers have a full portfolio of tools to do code testing. Project managers will find Requisite Pro – a requirements-management tool created by Rational – very easy to use. The bean counters even get something this time that shows the business process and ROI. As usual, IBM has focused on heavy code reuse, quality by design, and multiple group resources.

To get you going they'll send you a two-DVD set with all of the evaluation code. Then, over the next three months IBM will send you a multitude of artifacts and tutorials.

On a separate note, as an avid user of Google, I've noticed, as have many other people, that Google's sponsored links have become quite popular with companies looking for people to buy their stuff. For any of you who have not delved into Google's sponsored links, the program is very simple. You pick out the keywords you want your sponsored link to appear with when people search for the word or combination of words.

I recently typed in "IBM WebSphere" to see who was selling what and what I found surprised me. The very first ad has a headline that reads, "Get a Real AppServer. Cut licensing costs and train staff for less than WebSphere licensing." Google directs you to the JBossgroup.com Web site, so I clicked on it. The people over at JBoss may have a "real application server" but they seem to be missing the tools, documentation, and product extensions.

After careful reading I finally understood what they meant by "Get a Real AppServer." This is an open source kit, where you roll your own application server. I think that the open source movement is a good idea. But I don't think it's a good idea to imply that WebSphere is not a real application server. It just seems unfriendly.

My first thought was that maybe IBM had purchased the keyword JBoss so whenever people search for it they'll find a sponsored link that says, "Get a Real Appserver – one that has lots of tools, documentation, and product extensions."

Then I realized that by "Real Appserver" they meant JBoss. They want everyone to think "JBoss" is synonymous with "Real Appserver" so that if they decide they want one they will think of JBoss first. Maybe this type of advertising will work, but I doubt it. I have found that whenever someone starts off with how bad someone or something else is they usually have very little going for themselves.

The next link took me to shopping.com, where I found links to a whole lot of WebSphere software for sale at quite a few stores. I didn't realize that WebSphere had gone retail. The way it looked, maybe Wal-Mart will begin selling it soon. I was also happy to find our magazine for sale there.

From there I found VideoEd.com/b_websphere.html, a distributor of independent technology training that has interactive training courses for a huge amount of software, with WebSphere as one choice. I spoke with Bob Dumouchel, the owner – quite a friendly guy and very knowledgeable; he tells me he has 1,500 courses from 21 different producers with a strong focus on business and software skills. Medium-size businesses are their strike zone, but they work with a wide variety of companies. Much more important is that they will go out and find what you need if they don't have it.

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
MANAGING EDITOR JEAN CASSIDY
EDITOR NANCY VALENTINE
ASSOCIATE EDITORS JAMIE MATUSOW
JENNIFER VAN WINCKEL

WRITERS IN THIS ISSUE
Brad Bouldin, Kulvir Singh Bhogal, Jeff DeMent, Rahul Kitchlu, Sapna Mahwal, Jack Martin, Jim McQuaid, Tilak Mitra, Thina Natarajan, Guido Puglielli, Bob Sutor

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 RECEIVABLE CHARLOTTE LOPEZ
ACCOUNTS PAYABLE BETTY WHITE
ADVERTISING DIRECTOR ROBYN FORMA
DIRECTOR OF SALES AND MARKETING MEGAN MUSSA
ADVERTISING SALES MANAGERS ALISA CATALANO, CARRIE GEBERT
ASSOCIATE SALES MANAGERS KRISTIN KUHNLE, BETH JONES
CONFERENCE MANAGER LIN GOETZ
LEAD DESIGNER LOUIS F. CUFFARI
ART DIRECTOR ALEX BOTERO
ASSOCIATE ART DIRECTORS RICHARD SILVERBERG, TAMI BEATTY
WEB DESIGNERS STEPHEN KILMURRAY, CHRISTOPHER CROCE
CONTENT EDITOR LIN GOETZ
CIRCULATION SERVICE COORDINATORS SHELIA DICKERSON, EDNA EARLE RUSSELL, LINDA LIPTON
CUSTOMER RELATIONS MANAGER BRUNI STAROPOLI

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

NEWSSTAND DISTRIBUTION CONSULTANT:
GREGORY ASSOCIATES, W.R.D.S.
732 607-9941, BJGASSOCIATES@CS.COM

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

Architecture

Where Does SOA Fit?

BY BOB **SUTOR**

Bob Sutor, PhD, is IBM's director of WebSphere Infrastructure Software. He is responsible for business and product leadership for the WebSphere Application Server and WebSphere Studio software offerings. Previously, he was IBM's director of Web Services Technology, responsible for driving the cross-IBM Web services initiative to continue and advance IBM's leadership in providing Web services solutions, products, and services. Prior to this role, Bob led the global IBM-wide strategy effort for the development and promotion of key e-business industry open standards, including XML and Web services. He was a member of the OASIS Board of Directors and participated in various W3C working groups.

There's a lot going on in business and IT today, and it's rapidly becoming more important that you have a good understanding of e-business on demand, service-oriented architecture, Web services, and grid computing. I've been in more than one conversation and read more than one article where someone has defined these terms in, shall we say, "creative" ways. Let me try to put each in its place and show how they relate to each other.

I like to think that there are four major levels that describe how to successfully use IT to run your business. The top level is e-business on demand. This is all about business and describes the ideal way you want to operate to be more efficient, flexible, and responsive to market opportunities and your customers' needs. Of course, no business is perfect, but a good assessment of where you are now, where you want to be, and how you can get there will help you understand your chances of arriving at this nirvana. You'll know you're going in the right direction as it becomes easier to execute new business models and change old ones. More specifically, when you can quickly tune and modify your business processes and the ways they interact with your customers, partners, and suppliers, you will be operating in an on-demand way.

After all, isn't IT a major roadblock to really making this happen? What you need is an architecture that enables true distributed computing inside and outside your enterprise, that treats business applications as collections of services that can be shared or outsourced, and that uses standards to show how the pieces fit together and how they behave while operating. Of course, you want to do this in a secure, reliable, and manageable way. This is the second level – service-oriented architecture (SOA). It allows existing business software to continue to be used for critical processes, but extends its use for new interactions. When done correctly, it makes heterogeneous systems look homogeneous and highly leverages the infrastructure that you now have in place.

The third level is where we deal with specific technology. Web services standards describe the format of the information we use in business transactions. This includes the data (think of a purchase order or a request for a bank transfer), security information (encrypted and digitally signed patient records), and workflow (what gets done in what order). It also involves the additional information we need for reliable delivery of messages and performance management of the services. Web services technology was directly enabled by the widespread use of the Internet as well as the creation of XML. It is the concrete way of saying what we want done, what information we are using, and how we communicate this information to others.

A second piece of technology enabling SOA is grid computing. This is an important standards-based component of how we will deliver the qualities of service our customers will require if they are to remain our customers. For example, Schwab was able to reduce the time it took for an interactive customer portfolio analysis from roughly 4 minutes to less than 20 seconds using IBM-developed grid computing technology. Grid computing is about putting resources and information where they need to be as quickly as they need to be there. It lets you better use what you've already paid for and, perhaps, reduce your infrastructure by outsourcing resources. You pay only for what you need – on demand.

You need software and, at your discretion, services. This is the fourth and bottom level. If you are an IBM customer it means WebSphere software for transactions, Tivoli software for security and management, DB2 for data management, Lotus for collaboration, and Rational for the full life cycle of software development. Other vendors provide software as well, and there are open source options for some of the pieces. IBM and many other companies use the open source Eclipse Project as the basis for their tools. The most advanced software framework in the industry for building and using Web services and SOAs is the Java 2 Platform, Enterprise Edition, particularly the new version 1.4.

Service-oriented architecture is the glue that connects your business needs and models with the technologies that can make them a reality. There are more than 15 years of SOA research and development backing up our IT and business practices in support of the principles and goals of e-business on demand.

Deploying SOA in your organization can be done in an evolutionary manner, and clear patterns are emerging to make it a repeatable practice. It's time for you to seriously consider how SOA can help yours become a better business.

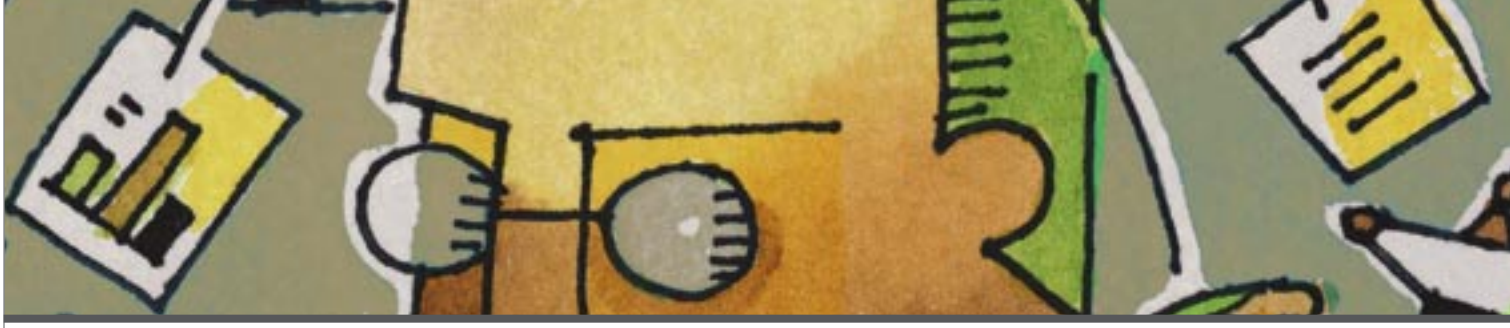

Going Vertical

IBM's move to an industry-specific strategy

INTERVIEW BY JACK **MARTIN**

WebSphere Journal editor-in-chief Jack Martin recently chatted with Paraic Sweeney, IBM's VP of Marketing for Industry Solutions and Business Integration, Software Group. In this exclusive interview, Sweeney discusses IBM's Middleware Industry Solutions initiative, the drivers behind adding an industry-vertical dimension to the business, and the company's decision in 1999 to focus on middleware and get out of the application business.

WebSphere Journal: Paraic, how do you spend a typical day?

Paraic Sweeney: My typical day – especially now that we are in a launch phase around the Middleware Industry Solutions initiative – is a combination of working with early-adopter customers and business partners who want to participate and join us in this new market approach that is an extension of our middleware business and communicating with our internal product and sales teams, who are all working to use this approach to the market as a mechanism to better address customer needs in the market. So it's a combination of customers, partners, and internal teams.

WJ: What type of customers are you actually dealing with? Which industries do they represent?

PS: We are actually working across a wide range of industries, given the breadth of IBM's market presence – and many of those are existing customers, but they may not necessarily be broad-based users of our middleware technology. We see customers in a number of industries who are working to respond to regulatory or industry mandates specific to their industry. For example, in the insurance area there are institutions that are figuring out ways to respond to the new federal mandates around HIPAA (Health Information Privacy and Accountability Act) compliance.

In consumer goods manufacturing, a major driver is the requirement that Wal-Mart is making across its network of suppliers to increase the efficiency of its supply chain. This includes the ability to trace items within Wal-Mart's supply chain using new technologies like RFID [radio frequency identification]. It's customers in a range of industries who are wrestling with these business issues and looking for the software assistance to work in concert with their existing infrastructure.

WJ: You're dealing primarily with the line-of-business people inside these companies?

PS: In many cases we will be partnering with systems integrators and consulting partners. Many of the topics I have just described are really at the confluence of business and IT. They're neither purely an IT discussion about technology, nor are they purely a business issue around packaging or pricing. It's really about how these two issues – in what up to now have been relatively distinct roles – have to interact with each other. And in many cases it would still be an IT-savvy line-of-business guy or a very line-of-business–focused IT guy, but it's at a relatively senior level that they are wrestling with these business decisions.

WJ: So the decisions are not made purely from an information technology perspective with these people?

PS: That's really at the heart of why we are adding an industry-vertical dimension to our business, which traditionally has had a technology orientation. In the past we've seen decisions made around "I want to standardize my database infrastructure for a more cost-effective way of storing and retrieving and aggregating data." It doesn't really necessarily have a specific business driver. However, we see spending on software increasing among the customers we're serving. It is shifting from purely an IT-focused decision and we are seeing additional projects with a clear business driver; therefore, the decision is no longer purely an IT one but is project based, and it is a collaboration of the sponsoring business unit and the IT function that supports it. They usually face quite aggressive implementation time schedules, which in many cases are externally mandated by legislation or a trading partner in their distribution or supply chain.

WJ: Can you give me just a flavor of what

type of business challenges these companies are looking at? What's a popular request that you hear?

PS: I think one challenge we are starting to see in a more broad-based way is how to implement RFID, not just to meet a customer requirement but to gain a business advantage in addition. It is a very new technology and there are significant players in the industry beginning to mandate its use. For example, Wal-Mart has asked its top 100 suppliers who ship goods into Wal-Mart's Texas distribution centers in 2005 to equip each pallet with an RFID tag. And now the federal government, specifically the military, has said "That's a pretty good idea – we've got a supply chain problem also. We want our top suppliers to do that also."

You can be fairly sure that once RFID proves its merit, it will be expanded from those distribution centers – from the top 100 to the next thousand and the next thousand after that. One company in that particular industry drives it and gets the business benefit out of it, and then you'd expect to see Staples and Target and other retailers adopting a similar approach in order to remain competitive in their industry. So, the question then becomes, "How do you go about it?"

If you look at an RFID implementation, it's a multidimensional problem in that you've got a lot of devices to manage; those readers will produce a lot of data. As with many new technologies, the standards on how to encode the information on the tags are changing, so you are going to have to transform and reformat the data.

Once you can automatically identify what goods are on the pallet, the retailer can now get a business value from this real-time data; for example, by doing an automatic three-way match. Wouldn't it be interesting before you actually accept a delivery, to do a check: "Well, what does the purchase order say and what does the manufacturer or distributor say that I am going to have on the pallet in terms of the advance ship notice? And then do a comparison across the three information sources – the purchase order, the advance ship notice, and the delivery data. We have described a piece of logic that does that match. If there is an exception, what do you want to do with that? It could be displayed for the procurement function and the supplier via a portal, which could in addition alert the store manager via pager or cell phone. What I have described is a classic information management, application integration, and portal scenario. So these questions – How do I respond to the RFID requirement? What technologies are available? How do I deploy? How do I manage it? – illustrate of the type of combined business and technology question that is being asked.

WJ: Taking RFID as an example, from what I've read and from what you are saying, it sounds as if you expect to see mass adoption of this technology over the next couple of years.

PS: I think over the next three to five years we will have a very significant adoption. What I have described is a starting point. You can obviously see it scaling in a number of ways: from one set of suppliers you can scale the number of suppliers. Instead of one set of distribution centers, RFID could be used for multiple distribution centers. The initial usage will put a tag on a pallet, but it could be all the cartons on a pallet. Potentially there is a massive adoption on the horizon.

WJ: Have you worked with any of the food supply people?

PS: On this particular issue, we are working with a number of consumer goods manufacturers.

WJ: How do you see this working in the automotive industry?

PS: There are similar requirements in the automotive world around early warning for warranties, and there are also some legislative pressures on auto manufacturers in response to issues with usage of parts from suppliers. These have been embodied in the TREAD (Transportation Recall Enhancement, Accountability, and Documentation) Act.

Another example is in connection with the maintenance of a car. There's now sufficient intelligence in an automobile to be able to detect problems early and communicate them back to the manufacturers so that they can prescribe appropriate maintenance. You're starting to see a variety of automatic capabilities being applied in a number of industries that deliver business benefit in addition to those changes that are being mandated by either market makers or regulators.

WJ: It sounds as if information technology is growing up to the next level, where business people come in with very, very clear ideas of what they are going to try to accomplish through the application of technology. Is that what you are saying?

PS: What I am saying is that there are a variety of business challenges that can be translated into the application of very specific technology, often with middleware combining existing applications and data. In many cases, the line of business that they're in is the driving force behind the initiation of these projects. So there is a growing awareness of what is possible – and the potential business benefit that can be obtained. This idea is the genesis or the reason behind expressing the value of our portfolio of middleware software in terms of the business issues that can be addressed. In addition to saying our application servers are J2EE compliant and our databases are scalable, our systems management manages a wide variety of devices, and our integration servers can support a wide range of B2B protocols, we are also starting to explain that we can assist with branch transformation in Banking, with dealer collaboration in Auto, with mobile service delivery in Telco, etc. That's the mechanism we're using to more directly address the issues our customers are working on but also, quite frankly, to allow our sales team to have a more meaningful conversation with people who are initiating software projects; that's the education program we are driving with our sales team. It's also a new focus for our engineering teams so that they think about these issues, and it's a partnering opportunity for systems integrators so we can work together to address these customer requirements.

WJ: You mentioned earlier that you are

also working with systems integrators and business partners. How do they dovetail into this initiative?

PS: I would describe the middleware solutions for each industry as having four components. The first part contains the core IBM middleware products, for example, DB2 database, Lotus Domino, Tivoli systems management, Rational tools, WebSphere Business Integration, etc.

The second part comprises our industry-specific middleware extensions that make our core products more relevant in that industry vertical – we have adapters to vertical ISV applications, process templates for common business processes, portlets for user roles, and I expect additional extensions in the future.

The next two main components to the solution are both partner based. Number three is the ISV application or applications that a customer has, because any solution has to work in the context of the customer's environment. The fourth is the industry know-how in terms of consulting and implementation services that can take those middleware products and deploy them in the application environment that the customer has, to solve the business problem. That know-how can be provided by IBM services – from IBM Global Services, obviously – and it can be provided by our network of business partners, global and regional systems integrators, and solution provider partners.

When you look inside many of the large global systems integrators – Deloitte, Accenture, and so on – they're largely structured around industry. Many of the more specialized systems integrators are even more deeply specialized around particular issues in an industry-vertical domain. When we talk to the integrators, they see a very natural fit between their industry know-how and our middleware software and the business opportunity for the combination. They are depending on IBM for two things. Number one is market awareness for our middleware software and the second thing is the technical support that allows them to successfully apply those technologies in their customer engagements.

WJ: Is IBM research involved with this initiative? If a customer has a really interesting high-value problem to solve, are you bringing in the Watson folks?

PS: In many cases, the Watson Research lab is involved in what they call the "first of a kind." We've got this program – as usual we have an acronym for it – called FOAK.

WJ: It's IBM. There has to be an acronym.

PS:: There was an article on this whole subject in *Information Week* just before the holidays. The cover is all about this topic and in the story is a discussion about a first-of-a-kind project with a customer called Boston Coach – a limo provider. IBM Research used some of the mathematical modeling that they have done for scheduling, which is typically used with airlines, where you have lots of flights coming in and out and so you need to optimize landing slots, planes, pilots, logistics, etc. We took that technology and that mathematical knowledge and applied it to Boston Coach's problem – scheduling limos in a city. Boston Coach is now able to expand their operations to other cities at a lower cost because they are able to share the systems that now automate a lot of that activity.

"Many of the topics I have just described are really at the confluence of business and IT. They're neither purely an IT discussion about technology, nor are they purely a business issue around packaging or pricing"

WJ: How are you getting the message out to the customers that all of these new vertical solutions are available? For instance, if I were an auto manufacturer, would I be coming to you or are you going out to that market?

PS: We are going out to the market – this article and others are part of an extensive communication initiative – but we are also training our customer teams – sales and product services – so they can initiate these discussions with our customers, and the partner community is a key channel.

WJ: So the partners in many cases have the relationship, and you are bringing in the technology muscle?

PS: Correct.

WJ: I've heard that you actually have an enormous portfolio of industries that you are covering. I have had a few people tell me that there are actually 12 of them. Is that true?

PS: It is, but let me qualify…. We are in startup mode – albeit starting very aggressively – and IBM covers 17 industries with a dedicated sales force. For this software initiative we are starting very aggressively in 12 of those

BREAK AWAY FROM THE PACK

kenetiks

WORLDWIDE JAVA TRAINING & MENTORING

www.kenetiks.com
888.KENETIKS

- ***WEBSPHERE TRAINING***
- ***PROJECT MENTORING***
- ***COURSE LEASING***

From development to deployment, Kenetiks shows you how to build J2EE applications for WebSphere Application Server and WebSphere Portal Server using WebSphere Studio Application Developer.

IBM Business Partner

17 and are working to provide a middleware-based solution for a set of the top business issues in those industries. We see that as an aggressive starting point, but I don't want to focus on the number 12 as being an absolute statement or an endpoint. It is the start of a journey that we believe will make us more relevant and establish more leadership within the middleware category.

WJ: Are you able to say which 12 you are starting with?

PS: Absolutely, the 12 are banking, financial markets, insurance, telecommunications, government, retail, consumer products, automotive, electronics, energy and utilities, life sciences, and health care.

WJ: If a business partner wanted to get involved with this program, what is the appropriate channel and mechanism for them to get involved?

PS: We largely categorize business partners into either solution providers and systems integrators, or application developers (ISVs). The current programs for these partners will be extended to accommodate this additional perspective.

WJ: If a customer wanted to find out about these vertical middleware offerings, how would they go about it? What should they do?

PS: As part of our sales training in January, we equipped the direct software sales teams with in-depth education on each of these middleware solutions, so your usual IBM account contact will be able to describe them and will be able to get you more information.

The IBM software account team that you would have a relationship with today around Tivoli, Rational, WebSphere, and so on will be familiar with these solutions. If you wanted to have a more detailed conversation to understand the technical components of it, to maybe develop a deployment plan within your shop, then they can bring in industry middleware solution specialists to assist.

Paraic Sweeney is the vice president of marketing, industry solutions, and business integration for IBM Software Group. He joined IBM in 1980 and has held a variety of international positions, including managing product development, financial planning, marketing, and consulting in Dublin; Paris; and Somers, New York. Paraic is responsible for marketing IBM's portfolio of middleware industry solutions that exploit their market-leading middleware infrastructure and open industry standards leadership. He was instrumental in the acquisition and integration of the CrossWorlds Software and Holosofx companies, both industry pioneers in process integration technology and now key components of the WebSphere product portfolio. Prior to this position he worked in a range of software leadership roles with IBM's Software Group including VP of Marketing, Internet Technologies, where he launched the WebSphere brand.

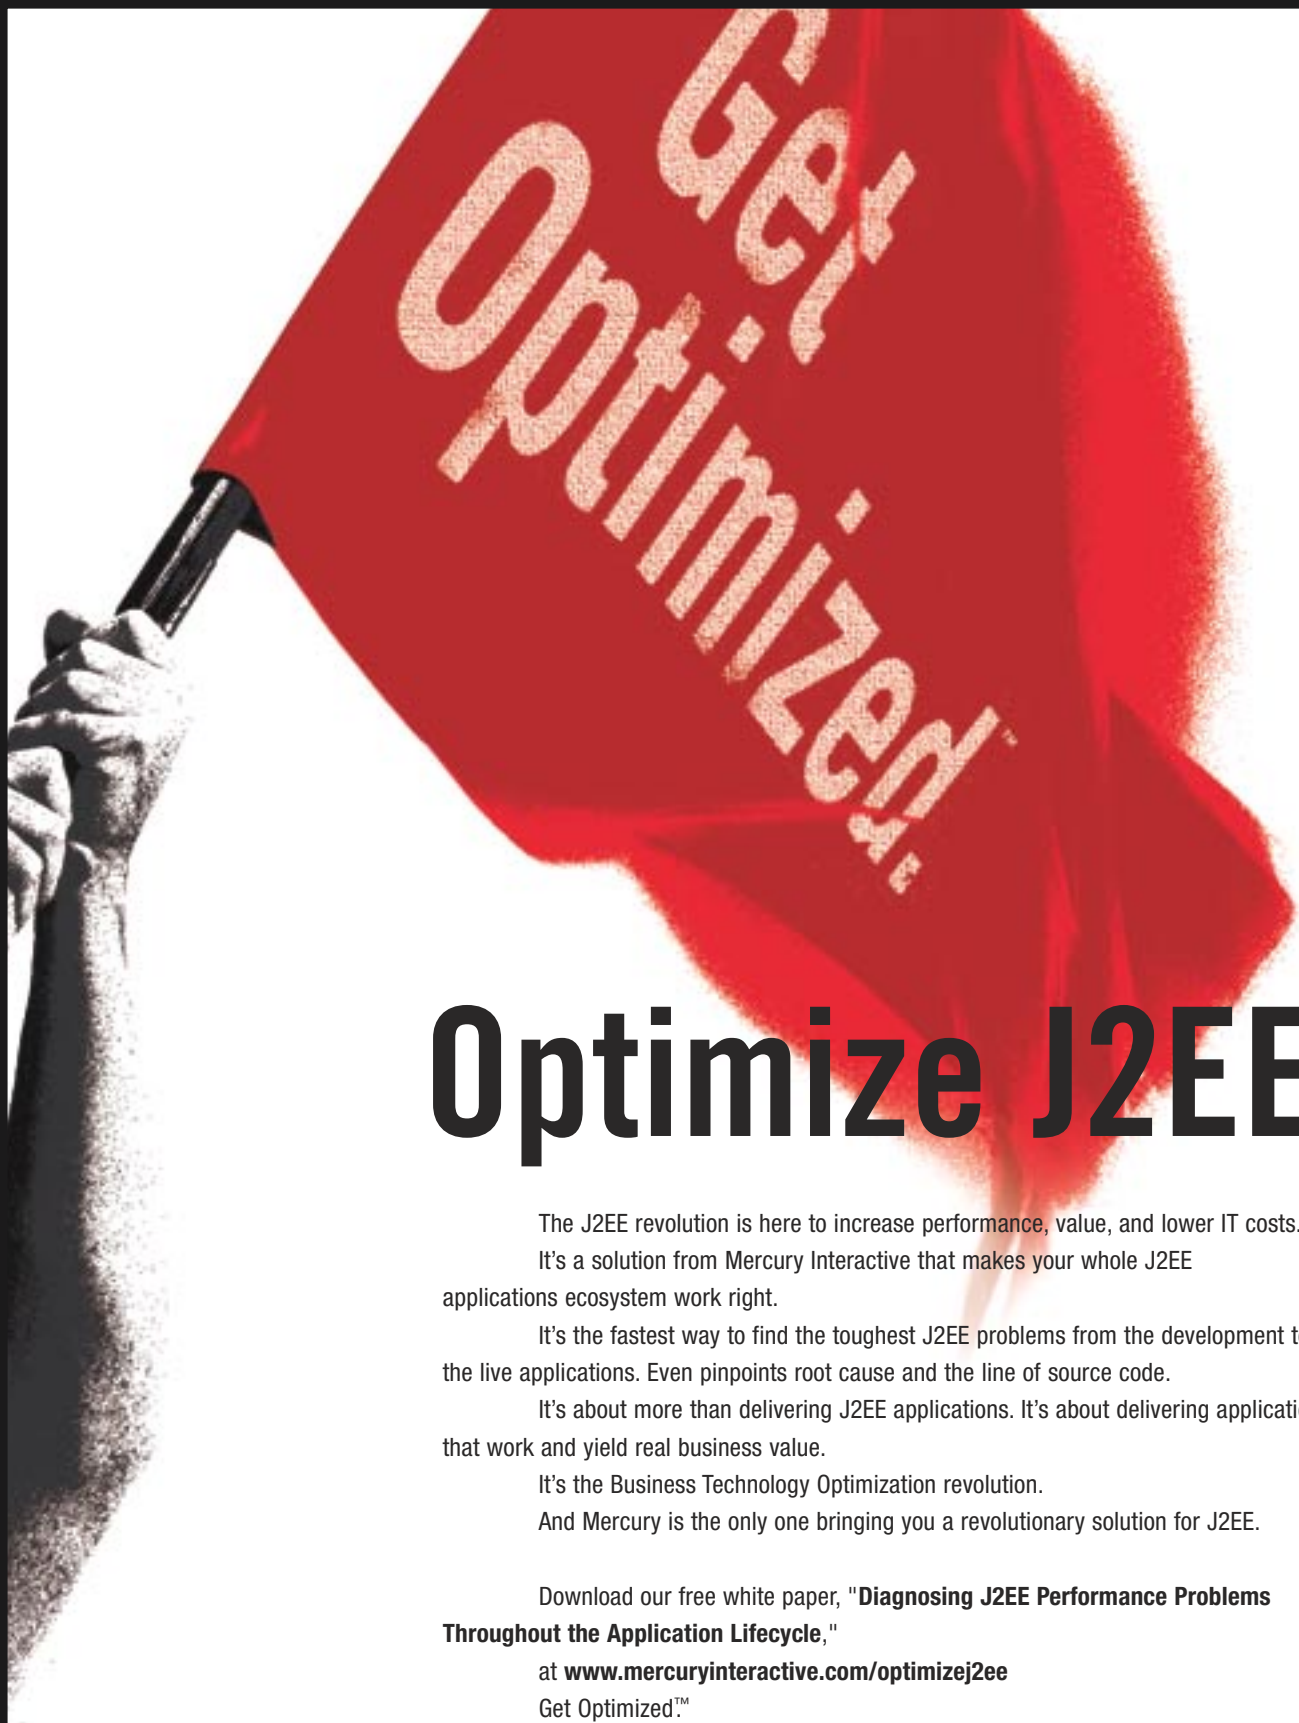

Optimize J2EE.

The J2EE revolution is here to increase performance, value, and lower IT costs.

It's a solution from Mercury Interactive that makes your whole J2EE applications ecosystem work right.

It's the fastest way to find the toughest J2EE problems from the development to the live applications. Even pinpoints root cause and the line of source code.

It's about more than delivering J2EE applications. It's about delivering applications that work and yield real business value.

It's the Business Technology Optimization revolution.

And Mercury is the only one bringing you a revolutionary solution for J2EE.

Download our free white paper, "**Diagnosing J2EE Performance Problems Throughout the Application Lifecycle**,"

at **www.mercuryinteractive.com/optimizej2ee**

Get Optimized.™

MERCURY INTERACTIVE

Accelerate each phase of the 'build, run, manage' life cycle

10 Benefits of Migrating to WebSphere Application Server v5.0

BY THINA NATARAJAN

As the gap increases between technology innovation and the constricted capabilities of inherited legacy systems, many organizations are evaluating the costs and benefits of migrating their infrastructures and enterprise applications to new generations of Web-based infrastructure technologies. Expanded abilities to accelerate the introduction of new services and reduce administrative overhead are important factors in the migration equation.

This article explores 10 fundamental benefits of migrating to IBM WebSphere Application Server (WAS) v5.0. Organizations can, for example, use WAS v5.0 to capitalize on open J2EE standards, as well as simplify development and deployment to expand their e-business capabilities. Migration also allows existing applications to harness new application server functionality and technology. A properly planned migration to WAS v5.0 can quickly yield return on investment (ROI) while providing organizations with the agility to quickly use IT to support new business services.

Benefit 1: Support Enterprise Java Open Standards

Today's enterprises typically include a number of disparate systems loosely linked by a series of vendor-specific and customized technologies. These systems and applications are unable to communicate efficiently with one another due to a lack of well-defined interfaces for application communication, and data interchange. Open standards such as J2EE are critical to standardizing and streamlining the ability of information to be shared among multiple environments.

J2EE 1.3 is an open, platform-independent, standards-based enterprise development platform. New J2EE 1.3 capabilities include additional open standards – such as Web services standards – to enhance scalability, flexibility, and manageability.

WAS v5.0 supports open standards via the J2EE 1.3 specification. Application developers, integrators, and designers using J2EE 1.3 are freed from the limitations imposed by a proprietary platform from a single vendor. A J2EE-based solution can be deployed on any J2EE-compliant application server. For example, organizations using J2EE can share information more easily with multiple departments and trading partners. The result is increased operational efficiency and a 360-degree view of customer activities across all lines of business.

Incidentally, the J2EE 1.4 specification was finalized in November 2003 and will be implemented in WAS v6.0. It is currently available as an early-access technology for developers.

Benefit 2: Ensure Forward Compatibility

Prior to J2EE 1.2, J2EE specifications did not ensure forward compatibility. The Java Community Process, the J2EE standards organization, has since ensured that all future J2EE updates will include migration capabilities. Future compatibility for J2EE is important because it enables organizations to protect and build on their technology investments without the concern that their IT infrastructures will become outdated and require a painful migration to a completely new J2EE environment.

Standardizing enterprise applications on a single version of J2EE (1.3 or later) will ensure forward compatibility and provide a technology foundation on which to build. Although such an initial migration may be resource intensive, subsequent updates will provide significant return on users' Java technology investments. Once the migration is complete, subsequent migrations are simply automated through the WAS v5.0 administrative console.

Benefit 3: Expand Development and Maintenance Capabilities

WAS v5.0 also provides simplified and expanded development and maintenance capabilities. Specifically, WebSphere Studio Application Developer v5.1 has replaced VisualAge for Java as the development tool for J2EE applications. Built on the open source Eclipse platform, WebSphere Studio enables open standards–based and scalable software development.

WebSphere Studio includes

built-in source control management through its support for team development tools, such as the open source Concurrent Versioning System (CVS) and IBM Rational ClearCase software configuration management application. Further, the model-driven development capabilities of WebSphere Studio include plug-ins for Rational XDE Professional that enable developers to make graphical models of their software. Graphical modeling tools can help minimize code rewriting during development via code generation, reverse engineering, and code-to-model synchronization capabilities.

Moreover, unlike VisualAge for Java, WebSphere Studio supports 100% J2EE development and deployment. WebSphere Studio provides a number of features that enable IT teams to accelerate development. Assembly tools and deployment descriptor editors enable development teams to create and modify deployment parameters and Enterprise Archives (EARs) from within the WebSphere Studio integrated development environment. WebSphere Studio also allows application debugging on remote application servers, which greatly accelerates development. It now also includes an agent controller and a Java profiling agent that enables application performance profiling through the Java Virtual Machine Profiling Interface (JVMPI).

The combination of WAS v5.0 and WebSphere Studio provides enhanced development and maintenance of next-generation applications. By tapping into the latest innovations provided by the J2EE 1.3 specification, WAS v5.0 provides greater control for application development, deployment, and performance.

Benefit 4: Streamline Deployment

Typically, the complexities of deployment have rested with application server administrators. By virtue of its support for J2EE, WAS v5.0 allows all application components to be packaged into a single EAR, which can be deployed on multiple WAS v5.0 instances. This capability allows organizations to realize the "write once, deploy anywhere" promise of J2EE.

Benefit 5: Strengthen Information Security

WAS v5.0 strengthens the management and implementation of organizations' information security safeguards by providing built-in support for authentication, authorization, and identity management. Authentication enables IT managers to confirm the identity of a user, typically by matching submitted username and password information with the user's stored credentials. After the authentication stage, each user is granted access to a specific set of files, databases, services, or administrative privileges based on an individual profile. Identity management software is used to automate and manage information security administrative tasks such as resetting user passwords. Enabling users to reset their own passwords, for example, reduces IT administration requirements while maintaining a high level of security.

WAS v5.0 authentication support includes the Java Authentication and Authorization Service and the Java Cryptography Extension. Additionally, Tivoli Access Manager (a.k.a. Tivoli Policy Director) provides a single source for sitewide authentication. Identity management standards supported in WAS v5.0 include Security Assertion Markup Language, Kerberos, and XML digital signatures. Further, WAS v5.0 supports Web services security specifications. The World Wide Web Consortium's WS-Security specification for Standard Object Access Protocol extensions enhances confidentiality and integrity for applications using Web services. WAS v5.0 allows for next-generation, build-to-integrate, and secure enterprise applications.

Benefit 6: Improve Productivity

WAS v5.0 – through built-in, programmable adapter technology – allows for rapid development and deployment of applications for Web services and other enterprise messaging. WAS v5.0 provides users with the ability to deploy Web services across multiple communication protocols, such as HTTP, JMS, or Java Remote Method Invocation/Internet Inter-Orb Protocol. Multiple adapters enable IT teams to realize the promise of Java's "write once, run anywhere" capabilities by enabling J2EE enterprise applications to seamlessly interconnect with myriad Java technologies and protocols. This enables organizations to work completely in the Java environment, thereby sidestepping enterprise infrastructure compatibility problems and virtually eliminating the requirement for custom coding.

WebSphere Studio development tools such as the Application Developer and the Enterprise Developer leverage the open source Eclipse framework and help integrate third-party packages and vendor tools. In turn, IT staffs are able to increase their productivity by using their preferred tools via Eclipse.

WAS v5.0 also supports content-based routing, load balancing, and distribution to achieve end-to-end optimization of applications. When combined with WebSphere Studio tools to help facilitate faster development, WAS v5.0 yields faster time-to-market and increased ROI.

Benefit 7: Gain Continued WebSphere Support

IBM ended support for WAS v3.5 in Q4 2002 and support for VisualAge for Java in 2003. Organizations that continue their operations on WAS v3.5 could face the additional risk of not having IBM expertise available to support their business operations. Migration to WAS v5.0 guarantees IBM support, along with compatibility with associated third-party tools and vendor packages. Migration also enhances connectivity to trading partners, customers, and vendors using the latest standards and technologies.

Benefit 8: Extend the Enterprise

Organizations that have the WAS v5.0 enterprise extensions will be

"Exposing existing processes to an enterprise is a challenging and difficult task"

better positioned to support service-oriented architectures (SOAs). Simply stated, SOAs enable organizations to link software components to create new business applications such as an application for processing customer orders. A key advantage of SOAs, however, is that developers are not required to spend inordinate amounts of time writing new code to create new business processes. Rather, developers use standard protocols such as Web services to connect components in order to build new applications.

All versions of WAS v5.0 (Express, Base, Network Deployment, and Enterprise) support Web services deployment. The WAS v5.0 Network Deployment and Enterprise versions also provide a Web services gateway, as well as a private Universal Description, Discovery, and Integration registry.

The support for Web services–based SOAs includes process choreography and application profiling. Exposing existing processes to an enterprise is a challenging and difficult task. Entering the enterprise arena following migration to WAS v5.0, by contrast, is a simple, automated process achieved through standard wizards and code generators.

Benefit 9: Flexible Management Capabilities

WAS v5.0 introduces a new administration model that includes a Web-based interface that administrators can manage from any Web browser – providing more flexible management capabilities. Neither a client applet install nor a dedicated machine is required for administration. Moreover, unlike earlier versions of WAS, all configurations are stored in XML files instead of proprietary repositories, which makes configuration transparent. The plain-text, self-describing XML configurations can be modified outside the WAS administration console and easily replicated to other WAS v5.0 installations.

Developers can also use additional tools to incorporate application and network management in the J2EE development environment via WAS v5.0 support for Java Management Extensions (JMX). An optional extension to the standard Java Developer Kit, JMX can be used in place of SNMP. A key benefit is JMX's ability to facilitate the centralized management of J2EE objects and provide added structure to include performance management early in the development process.

WAS v5.0 also provides the wsadmin scripting framework for easy access to server administration features. Based on the open source Bean Scripting Framework, WebSphere Application Server's wsadmin feature allows for the automation of many daily administrative tasks, which frees administrators to focus on other activities such as application and deployment optimization.

Benefit 10: Enhance Reliability, Scalability, and Availability

WAS v5.0 uses cluster and workload dynamics to provide a high level of scalability, availability, and reliability. Its components and configurations will support the ability to deploy high-performance Web environments. To eliminate a single point of failure, organizations can facilitate clustering by using the dispatcher and content-based routing components, multitier load balancing, and Edge Side Includes caching. Further, WAS v5.0 includes a security server and a name server in every application server process to increase the protection of information.

Conclusion

Organizations are continually driving their IT departments to extend their capabilities in order to support expanded business requirements. IT departments, however, are increasingly challenged to enhance e-business capabilities with fewer resources.

Given IT's critical link to supporting business operations, it is essential that organizations use infrastructures that integrate with existing systems yet have the agility to support new technologies. WAS v5.0 provides organizations with the open J2EE capabilities and myriad enhancements that enable IT teams to accelerate each phase of the "build, run, manage" life cycle.

Resources

- Beaton, W. (2003). *Migrating to IBM WebSphere Application Server, Part I: Designing Software for Change:* www-900.ibm.com/websphere/techjournal/0107_beaton/index_eng.shtml
- Chalmers, D., et al. (2003). *A Quick Guide for Migrating to WebSphere Application Server V5:* www-900.ibm.com/developerWorks/cn/wsdd/library/techarticles/0304_chalmers/chalmers_eng.shtml
- Eeles, P., et al. (2003). *Building J2EE Applications with the Rational Unified Process.* Addison-Wesley Pearson.
- Wahli, U., et al. (2003). *WebSphere Studio Application Developer Version 5 Programming Guide:* http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246957.html?Open
- Yu, S., et al. (2003). *Migrating to WebSphere V5.0, An End-to-End Migration Guide:* http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246910.html?Open

Thina Natarajan is a principal solution architect for the Application Infrastructure Management (AIM) group at Candle Corp. Thina optimizes IT architecture, development, and deployment phases for global clients. He has recently been focusing on J2EE performance tuning best practices. thina_natarajan@candle.com

Using JAXB to Develop Enterprise Applications

JAXB offers power, ease, and flexibility

BY TILAK **MITRA**

JAXB (Java Architecture for XML Binding) provides a convenient way to bind an XML Schema to a representation in Java code and makes it easy for developers to incorporate XML data and processing functions into applications based on Java technology without having to know much about XML itself.

It has been proven over the past few years that the best form of information exchange (in a typical B2B or B2C scenario) is through XML. There are various XML-based standards (schemas) for both the horizontal and vertical market sectors, and there are ongoing efforts to move toward a standardized format.

With the proliferation of XML-based information exchange, the industry will write lots of Java code to consume XML Schema–based documents. JAXB (Java Architecture for XML Binding) provides a convenient way to bind an XML Schema to a representation in Java code. This makes it easy for developers to incorporate XML data and processing functions in applications based on Java technology without having to know much about XML itself.

This article will show developers how to utilize the power, ease, and flexibility of JAXB from within the WebSphere Studio Application Developer (referred to as WebSphere Studio), while developing real-world enterprise applications. With JAXB, developers no longer need to learn the intricacies of XML parsing techniques!

What You Need to Know

Although developers need not worry about XML parsing techniques anymore, they may need to do some reading up.

JAXB is built upon XML Schema, so a sound understanding of XML Schema is imperative in order to harness the power of JAXB. Business domain objects and their structural relationships can be well represented through an XML Schema, provided you have a good understanding of its power. The essence of JAXB lies in its customization of the XML Schema for Java objects. These customizations follow the rules of XML Schema. Adding JAXB customizations to a business domain object model (represented in XML Schema) is simple once the hard work of creating the schema is accomplished.

This article assumes you understand:

- XML Schemas (the more detailed the better)
- JAXB
- WebSphere Studio 5.1

Software and Packages Used

You must have the following products to complete the steps in this article:

- WebSphere Studio Application Developer 5.1
- Java Web Services Developer's Pack (JWSDP v1.3)

While we are all familiar with WebSphere Studio, JWSDP requires a little more introduction.

JWSDP from Sun Microsystems is a free integrated toolkit you can use to build, test, and deploy XML applications, Web services, and Web applications with the latest Web services technologies and standards implementations. JAXB is an essential part of this package. JWSDP 1.3 must be downloaded and installed (see the Development Environment Setup section for a detailed explanation of the process).

JAXB Basics

In object-oriented programming (OOP) there are two pivotal concepts: classes and objects. *Classes* provide the structure for software concepts or entities, whereas *objects* are live instances of the classes. A similar analogy can be used in XML representation. An XML Schema can be conceptualized as the allowable structure and constructs that can be used in the creation of an XML document (which conforms to the schema). The schema (in XML) and classes (in OOP) are the conceptual building blocks, whereas documents (in XML) and objects (in OOP) are live instances that conform to their respective conceptual building blocks.

Working with Java objects and classes is fundamentally different from working with XML. JAXB brings in the idea of data binding, which creates a correspondence between the XML schemas and the Java classes, and then

utilizes the mapping to convert XML documents to and from Java classes. The JAXB schema compiler creates Java classes and interfaces based on the structure of the XML Schema. JAXB libraries are used in marshaling and unmarshaling. *Marshaling* is the process of turning one or more Java objects into an XML document, whereas *unmarshaling* is the reverse process – creating a Java object from an XML document.

To use JAXB in a Java application the first step is to run the JAXB compiler (XJC) to create the Java classes and interfaces. For every element in the XML Schema, one or two Java interfaces and corresponding implementation classes (which implement the interfaces) are generated. The implementation classes are generated in a package that is separate from the one in which the interfaces are generated. Other than the interfaces and the implementation classes, JAXB-specific classes are generated to perform marshaling and unmarshaling, and to create instances of the implementation classes, among other things. (*Note:* The implementation classes are only instantiated through a factory class generated by the JAXB compiler). The two classes that are primarily used in working with JAXB-created Java objects are:

- ***JAXBContext:*** Used for marshaling, unmarshaling, and validating XML documents
- ***ObjectFactory:*** Used to instantiate the various implementation classes

Development Environment Setup

The first step is to download and install the JWSDP. As a part of JWSDP, the JAXB compile and runtime libraries and executables are installed. The next step is to download the JAXBTestEAR.zip file from www.sys-con.com/websphere/sourcec.cfm and extract the contents into the C:\temp folder. Next, create a new workspace and import the JAXBTestEAR file:

FIG 1: **EAR FILE TYPES**

1. Create a folder called JAXB and another called workspace under JAXB in the C:\ drive.
2. Open WebSphere Studio and point the workspace to C:\JAXB\workspace.
3. Select File>Import from the WebSphere Studio menu.
4. Select the EAR file type to import (see Figure 1).
5. Browse the file system and select the JAXBTestEAR.ear file. Click Finish.
6. The imported project in the J2EE Perspectives, Project Navigator tab should look like Figure 2.

With the initial setup complete, the JAXB compiler needs to be set up as an External Tool through the WebSphere Studio menu.

1. Select External Tools from the WebSphere Studio menu as shown in Figure 3.
2. Click New to create a new configuration. In the Main tab, key in the Name, Location, Working Directory, and Arguments as shown in Figure 4. (Browse the File system and Workspace in order to obtain the values.)
3. Switch to the Resource tab and check the checkboxes. Highlight the $resource Scope Variable and then choose JAXBTest as the specific resource.
4. Click Apply to save the configurations. Click Run to run the JAXB compiler (xjc.bat) and generate the Java classes and interfaces. The output in the output console will resemble Figure 5.

This completes the setup of the development environment inside WebSphere Studio.

Example Scenario

The example we will develop is a simple GUI menuing system. A widget is a conceptual element that defines the whole set of drawing components that constitute a graphical user interface. Examples of widgets are rectangles, squares, circles, and any other drawable visual component. Our simplistic GUI consists of three visual components in the form of rectangles, squares, and circles. The Widgets component encapsulates a list of the visual components (rectangles, squares, and circles) in any random order. A client can construct the visual components from a given XML document and then retrieve the details of each visual component (to be used in any way the client requires).

FIG 2: **IMPORTED PROJECT ON THE PROJECT NAVIGATOR TAB**

Listing 1 illustrates the XML Schema (for the example scenario) in some detail. The xs:annotation element is a container for the xs:appinfo and xs:documentation elements, which contain additional information. These two elements are dedicated to holding machine-processable (xs:appinfo) information and human-readable documentation (xs:documentation).

The section highlighted in blue denotes the JAXB-specific global custom-

izations of the schema that are applicable to the entire schema file. The section highlighted in red denotes the JAXB-specific customizations in which the package where the generated files will be placed is defined (in this case, com.ibm.domainobjects).

In the globalBindings section the collectionType = java.util.ArrayList is highlighted in blue. It denotes that any Collections that are created inside the Java objects all conform to the List interface in Java and are actual instances of ArrayList.

The xjc:serializable element makes sure that all the generated classes are serializable. (This is particularly important when objects are sent and received over the wire while data transfer between various application tiers is performed.)

A superClass called Shape is also defined. This denotes that all the elements defined in this schema have a common superclass called Shape. I show this here in order to illustrate that the classes that are created from the schema can also refer to external classes, i.e., classes that are not created by the JAXB compiler.

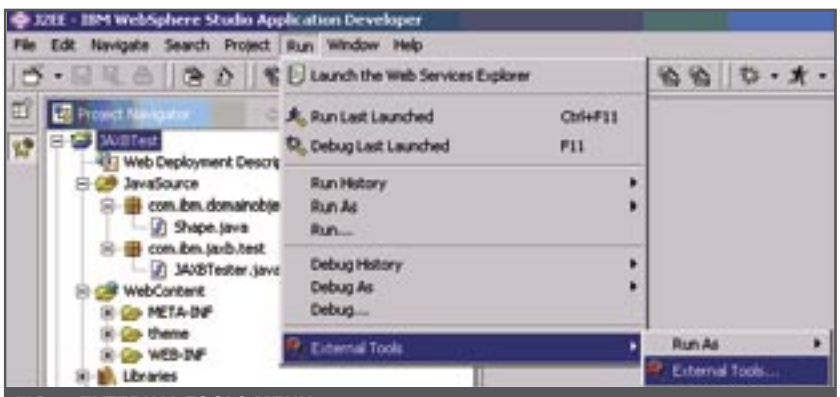

FIG 3: **EXTERNAL TOOLS MENU**

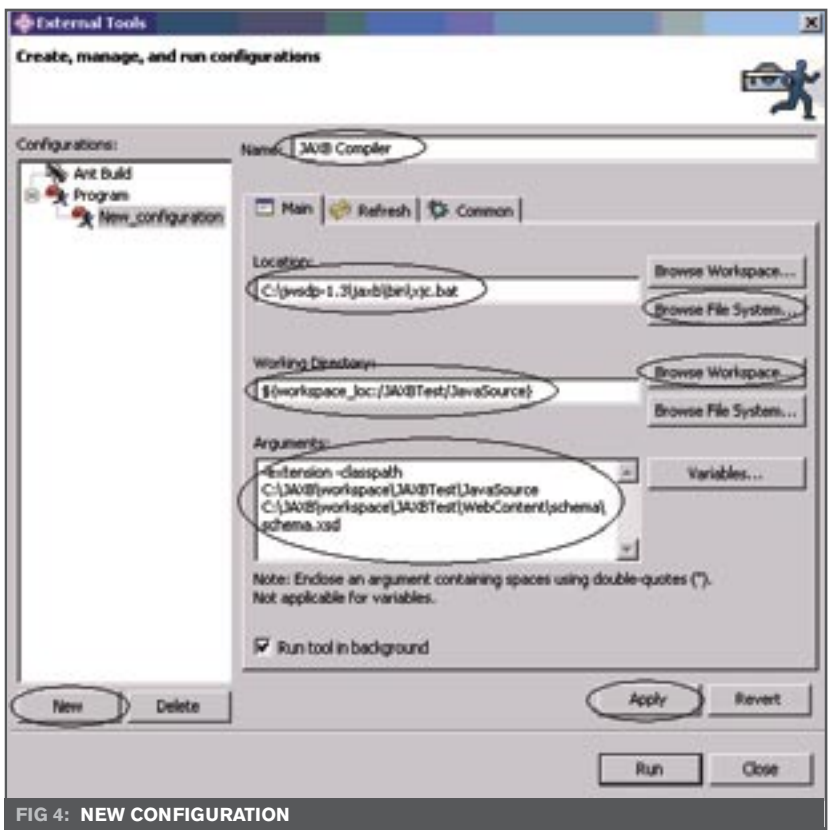

FIG 4: **NEW CONFIGURATION**

The definition of the Widgets element shows that there is a property called "Widgets" that is defined. Had this property not been defined, the binding compiler would have autogenerated a name for the list of choices (for the rectangles, squares, and circles). The autogeneration usually takes the names of the subelements and combines them with "Or". Hence the accessor name for the list would have been something like:

```
java.util.List getRectangleOrSquareOrCircle();
```

Clearly this is not meaningful. The jxb:property name="Widgets" denotes that the accessor method will be getWidgets, which is clearly much more intuitive than the autogenerated name.

In the example, I will demonstrate:

1. How to set up the usage of the external JAXB compiler from WebSphere Studio
2. How to use the JAXB compiler to generate Java interfaces and classes
3. How to create the Java classes from an XML document
4. How to instantiate Java classes, set their attributes, and then generate the corresponding XML document (the reverse of the previous step)

Various other customizing facets can be used in an XML Schema. The detailed specification for JAXB Customizations is found in the Resources section in the links for the JAXB Specification, Customizing JAXB Schemas, and Using JAXB.

Using the JAXB Generated Classes

The JAXB compiler is run only once. There is no need to run it again if the schema is not changed. However, if the schema undergoes any changes, then the JAXB compiler preprocessing step must be executed again in order to regenerate the Java interfaces and classes. In this event, it is advisable to delete all of the generated classes from a previous run of the JAXB compiler before running this step again. This is particularly helpful in cases in which elements in the schema are deleted. The generated classes from the new run of the compiler do not delete the old generated classes and hence some unused classes from the previous version of the schema will be left behind.

Once all the classes are generated, the fun begins, with the bulk of the tedious (if you consider it so) setup work already completed. It is time to use the generated classes in order to convert XML documents into Java classes and vice versa. We will concentrate on a client class called JAXBTester. This class demonstrates two activities. It first reads in a given XML document from which it creates the Java object structure/tree.

Analysis of JAXBTester

JAXBTester has three methods. A brief explanation of each of the methods is in order.

- ***createContext:*** Creates an instance of the JAXBContext class from which we can get a handle to an instance of the Unmarshaller and Marshaller instances.
- ***unmarshallIt:*** Accepts an XML document (which in

MQSoftware
SEE IT WORK.™

8,557 orders

flowed through your system flawlessly today.

1 didn't.

Unfortunately, one failed transaction a day

will cost you **$4.3 million** this year.

– If only you saw it sooner.

.:'':..:.:: Real-time monitoring of all your critical business transactions.
:.'':.::.:.::'. Because finding out tomorrow might be too late.

MQSoftware
Q Nami!™

IBM Business Partner

To learn more about Q Nami! and the benefits it brings to your business, register to attend a free online Web seminar.
Visit **www.mqsoftware.com/qnami** for seminar dates, or call **800.998.2858** for details.

our case is the "schema.xml" file found in the schema folder under WebContent). It creates an instance of the UIWidgets class (the root element in the XML document), then invokes the getWidgets accessor method on the UIWidgets instance and iterates through the list of contained Widgets, retrieving the various Widget-specific attributes in each iteration.

- ***marshallIt:*** Illustrates how we can instantiate any element in the XML Schema (for which there are generated Java objects), set the various attribute values, and then add each object to its container. An instance of ObjectFactory is used to create the various instances of the Java objects, for example RectangleImpl, CircleImpl, etc. The getWidgets method on the instance of UIWidgetsImpl returns a handle to an instance of a live List to which the child elements (instances of RectangleImpl, SquareImpl, and CircleImpl) can be added. Once the object structure is created, a single method call on the instance of the Marshaller generates the XML representation of the object structure.

Notice how not a single line of XML-specific code needs to be coded by the developer; the specifics of the conversion are all contained inside the generated classes.

Running the Sample

Before running the client class several steps must be followed in sequence:

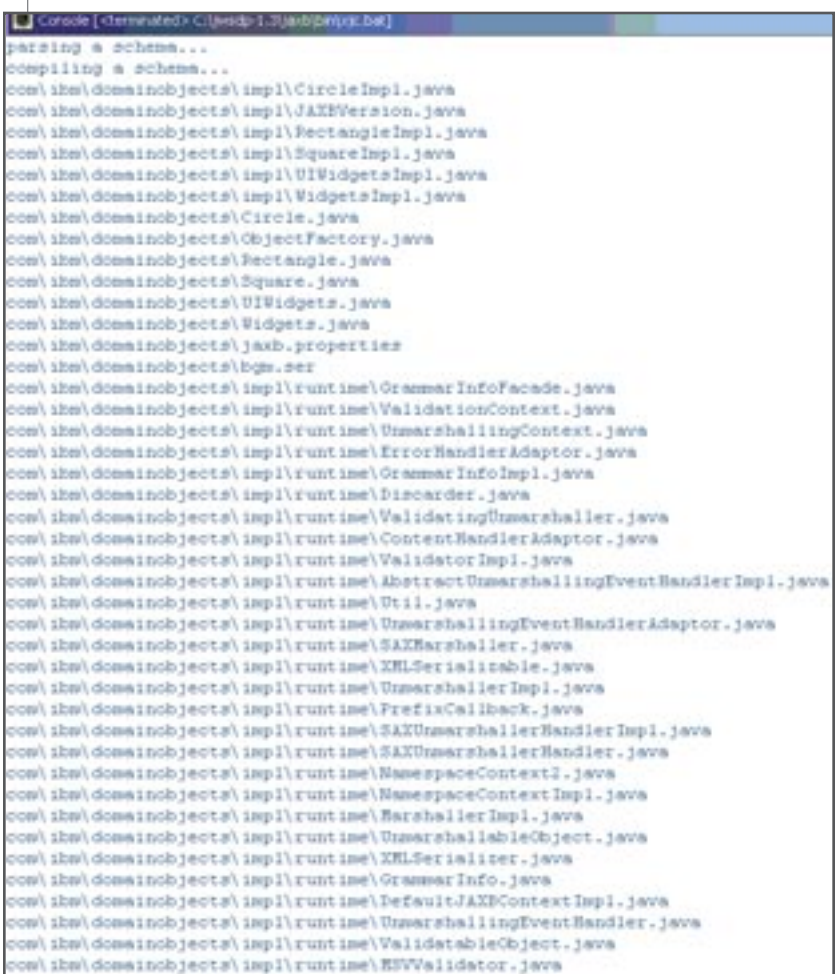

FIG 5: JAVA CLASSES AND INTERFACES

"Not a single line of XML-specific code needs to be coded by the developer; the specifics of the conversion are all contained inside the generated classes"

1. The JAXB compiler must be run to generate the Java object tree. (This step should already have been completed.)
2. The import statements that are commented out in the JAXBTester must be uncommented. The method bodies of unmarshallIt and marshallIt that are commented out need to be uncommented prior to compilation of JAXBTester. These portions are commented out because the classes that are referenced (both in the import statements and inside the methods) are all generated classes. When the EAR file was imported, these generated classes did not exist; hence, prior to Step 1 JAXBTester would not have compiled.
3. The value of the absPathToXML, which denotes the path to the schema.xml file, must be changed (if required) to reflect the path to the file in the file system.

Now you're all set to run the sample and see the results for yourself.

Conclusion

This article introduced you to the basics of JAXB. The most important aspect of this article was the demonstration of how JAXB can be used in a J2EE application and how the all development can be achieved from our friendly and most loved IDE – WebSphere Studio.

The best way to learn JAXB is to get very familiar and comfortable with XML Schema and then learn the tricks of JAXB customizations. The best way to learn XML Schema is to read a book on it, and then apply the concepts in real-world scenarios, exploiting the various features. (The O'Reilly book *XML Schema* is my favorite.)

Resources

- *The Java Web Services Developer Pack 1.3:* http://java.sun.com/webservices/downloads/webservicespack.html
- *Data Binding with JAXB:* www-106.ibm.com/developerworks/xml/edu/x-dw-xjaxb-i.html
- *The JAXB Specification:* http://java.sun.com/xml/downloads/jaxb.html
- *Customizing JAXB Schemas:* http://java.sun.com/xml/jaxb/users-guide/jaxb-custom.html
- *Using JAXB:* http://java.sun.com/xml/jaxb/users-guide/jaxb-using.html
- *Java Architecture for XML Binding (JAXB):* http://java.sun.com/developer/technicalArticles/WebServices/jaxb/index.html

Tilak Mitra is a certified IT architect in AIS, EAD Southeast sector, in the Fort Lauderdale/Miami area. He specializes in mid- to large-range enterprise and application architectures based on J2EE, MQ, and other EAI technologies. tmitra@us.ibm.com

LISTING 1

```
<?xml version="1.0" encoding="UTF-8> <xsd:schema xmlns:
xsd="http://www.w3.org/2001/XMLSchema"
         xmlns:jxb="http://java.sun.com/xml/ns/jaxb"
         jxb:version="1.0"
         xmlns:domainObjects="http://www.ibm.com/
         domainobjects"
         targetNamespace="http://www.ibm.com/domainobjects"
         xmlns:xjc="http://java.sun.com/xml/ns/jaxb/xjc"
         jxb:extensionBindingPrefixes="xjc">

    <xsd:annotation>
      <xsd:appinfo>
        <jxb:globalBindings collectionType="java.util.
        ArrayList"
                 fixedAttributeAsConstantProperty="true"
                 generateIsSetMethod="false"
                 enableFailFastCheck="false"
                 choiceContentProperty="false"
                 underscoreBinding="asWordSeparator"
                 typesafeEnumBase="xsd:NCName"
                 typesafeEnumMemberName="generateError"
                 enableJavaNamingConventions="true"
                 bindingStyle="elementBinding">
        <xjc:serializable />
        <xjc:superClass name="com.ibm.domainobjects.
        Shape" />
        </jxb:globalBindings>

        <jxb:schemaBindings>
            <jxb:package name="com.ibm.domainobjects" />
        </jxb:schemaBindings>
      </xsd:appinfo>
    </xsd:annotation>

    <xsd:element name="UIWidgets" type="domainObjects:
     Widgets" />
    <xsd:complexType name="Rectangle">
        <xsd:attribute name="width" type="xsd:integer"/>
        <xsd:attribute name="height" type="xsd:integer"/>
    </xsd:complexType>
    <xsd:complexType name="Square">
        <xsd:attribute name="length" type="xsd:integer"/>
    </xsd:complexType>

    <xsd:complexType name="Circle">
        <xsd:attribute name="radius" type="xsd:integer"/>
    </xsd:complexType>

    <xsd:complexType name="Widgets">
        <xsd:choice minOccurs="0" maxOccurs="unbounded">
            <xsd:annotation>
                <xsd:appinfo>
                    <jxb:property name="Widgets" />
                </xsd:appinfo>
            </xsd:annotation>
        <xsd:element name="rectangle" type="domainObjects:
        Rectangle" />
        <xsd:element name="square" type="domainObjects:
        Square" />
        <xsd:element name="circle" type="domainObjects:
        Circle" />
        </xsd:choice>
    </xsd:complexType>
</xsd:schema>
```

WebSphere Journal Advertiser Index...

COMPANY	URL	PHONE	PG
BOWSTREET	WWW.BOWSTREET.COM/EASY	978-863-1500	17
CYANEA	WWW.CYANEA.COM/WSDJ/NOMOREMADNESS/HTML	877-CYANEA8	7
H&W COMPUTER SYSTEMS	WWW.HWCS.COM/WSJ1.ASP	800-338-6692	3
INFORMATION STORAGE + SECURITY JOURNAL	WWW.ISSJOURNAL.COM	888-303-5282	33
IT SOLUTIONS GUIDE	WWW.SYS-CON.COM/IT	888-303-5282	37
JAVA DEVELOPER'S JOURNAL	WWW.SYS-CON.COM/JAVA	888-303-5282	45
KENETIKS	WWW.KENETIKS.COM	888-KENETIKS	11
MERCURY INTERACTIVE	WWW.MERCURYINTERACTIVE.COM/OPTIMIZEJ2EE	800-837-8911	13
MQSOFTWARE	WWW.MQSOFTWARE.COM/QNAMI	800-998-2858	21
NETIQ	WWW.NETIQ.COM/SOLUTIONS/WEB	408-856-3000	52
PROLIFICS	WWW.PROLIFICS.COM/WEBSERVICES	800-675-5419	2
QUEST SOFTWARE	HTTP://JAVA.QUEST.COM/JPROBE.WDJ	800-663-4723	51
VERSATA	WWW.VERSATA.COM/BUSINESSLOGICDESIGNER	510-238-4100	29
WILY TECHNOLOGY	WWW.WILYTECH.COM	888-GET-WILY	5

WebSphere Journal Coming Next Month...

INTEGRATION

HATS Off to Legacy

BY LOUIS J IACONA

PERSONALIZATION

Getting Personal

BY BHADRI MADAPUSI

INDUSTRY INSIGHT

Five Steps to Successful Strategic Partnering

BY BUELL DUNCAN

ILLUMINATING WEBSPHERE

Hunting Java Memory Leaks

BY LLOYD HAGEMO & RAVI KALIDINDI

Expediting content delivery with the dynamic cache service

Caching In

BY BRAD BOULDIN
& KULVIR SINGH BHOGAL

The quest for increased application performance is a science in itself. IBM WebSphere Application Server includes a powerful caching technology called the dynamic cache service, which you can employ in your Web applications to dramatically improve performance.

In this article, we'll use the dynamic cache service to increase the performance of a simple Web application. We will then use Apache JMeter, an open source load generator, to load-test our cached application in order to measure the application performance gains achieved by leveraging the dynamic cache service.

As you will see, the potential performance gains realized by using the dynamic cache service can be achieved with minimal effort. No changes to your application code are required, which means there is no impact on application portability.

The dynamic cache service is the name given to a set of services provided by WebSphere Application Server. The types of caching services include servlet and JSP caching, a Java object cache facility, a cache for Web services, and a command cache. In this introductory article our example will focus on the use of the servlet and JSP caching facility.

Enter Stage Left – Our Application to Tweak

Our demonstration application is implemented as a servlet that fetches a large piece of content from a remote HTTP server, then performs some text processing and returns the result to the end user. Our servlet grabs content from the Gutenberg Web site, which hosts downloadable books in ASCII format (www.gutenberg.net). For those unfamiliar with the Gutenberg project, this site hosts downloadable books that have expired copyrights and are therefore freely available.

Our example servlet will perform some text processing on the downloaded book: it will italicize the introductory section in the book's text that describes the Gutenberg project. We will test the performance before and after turning on caching and then compare the results. Once caching is enabled, the book will be downloaded on the first servlet request and presented to the user. Subsequent servlet requests will be served from the cache.

The servlet code (see Listing 1) is simple and straightforward, and uses standard Java packages and classes to read, parse, and transmit the content.

The servlet opens a URL connection to a book on gutenberg.net (see Figure 1) and copies the output to the servlet's output stream. The first section of the book contains the Gutenberg preamble (disclaimers and such), which the servlet puts in italics. Once the servlet reaches the end of the preamble (marked in the text stream with "*END*" at the end of a line), the servlet stops italicizing and switches to a more readable font for books (Times). The servlet also logs how long it took to produce the result, which has the side effect of showing us when the servlet is called (versus served from the cache). This will be handy once we enable caching because the log will show that WebSphere is using the cached result and bypassing the servlet call completely.

You can download an EAR file containing the servlet and test the performance on your own system. The source code for this article can be downloaded from www.sys-con.com/websphere/sourcec.cfm.

To see the effect of using dynamic caching in our application, we'll juxtapose the performance of our application with and without caching enabled. To study our application's performance, we'll leverage Apache JMeter. It is beyond the scope of this article to delve into the details of how to use JMeter, but you can learn more about how to use JMeter at www.devx.com/webdev/Article/17950/0. JMeter allows us to simulate a load of requests against our application server and measure subsequent performance.

We used JMeter v1.9.1 from http://jakarta.apache.org/jmeter. Our test environment consisted of one machine running WebSphere Studio Advanced Edition 5.0.1 with the WebSphere v5.0 Test Environment on a Windows 2000 Pro machine configured with a single AMD 3000+ processor and 1GB of RAM. We ran JMeter on the same machine. Since this example is "I/O-bound," running JMeter on another machine does not significantly affect performance. We used WebSphere Application Server

5.0.1, but other versions will yield similar results.

We configured JMeter with a "Jakarta Users" thread group, using 5 threads and 100 iterations (see Figure 2).

Under the thread group we added an HTTP request to access our servlet. We used port 9080 to access our servlet's Web container directly (see Figure 3). Figure 4 shows the results of our baseline (precache) test.

It's not too surprising that the average servlet response time is almost 15 seconds; it takes time to process and transmit 901K of text. Because each request takes so long (by design), our throughput is a mere 18.7 requests per minute. Our next step is to enable the dynamic cache service and measure the performance improvement.

Enabling the Dynamic Cache Service in WebSphere Application Server

Enable the dynamic cache service in the WebSphere administrative console as follows: open the administrative console, click Servers > Application Servers in the navigation tree, click on your server, and select Dynamic Cache Service under Additional Properties. Select Enable service at server startup in the Startup state field and click OK.

Enable servlet caching as follows: in the administrative console, click Servers > Application Servers in the navigation tree, select your server, and click Web Container. Check the Enable servlet caching checkbox under the Configuration tab. Click Apply or OK.

To complete the configuration changes, click Save on the top menu bar, click the Save button, and then restart the application server.

Configuring Our Servlet to Use Dynamic Cache

To specify the cache settings for our servlet, we created a cachespec.xml file (see Listing 2), located in the WEB-INF subdirectory of our WAR file. Let's take some time to discuss some of the cache settings in this file.

- ***sharing-policy:*** We set this to shared-pull, which means that the cached data will be shared across nodes in a clustered environment, and that each node will have to pull (request) the data from other nodes. The other options are not-shared, shared-push, shared-pull, and a

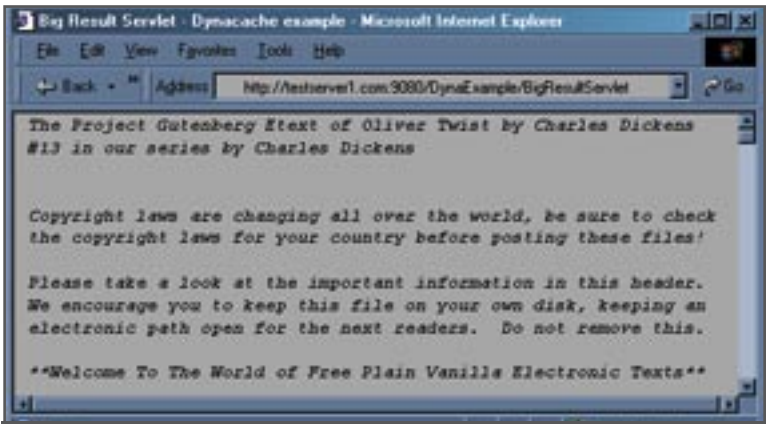

FIG 1: THE SERVLET OPENS A URL CONNECTION TO A BOOK ON GUTENBERG.NET

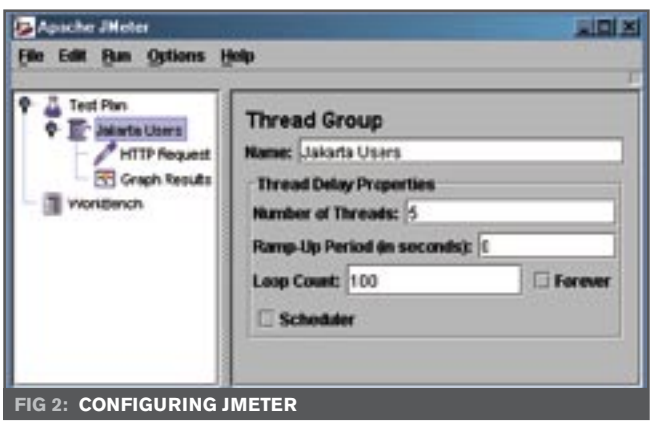

FIG 2: CONFIGURING JMETER

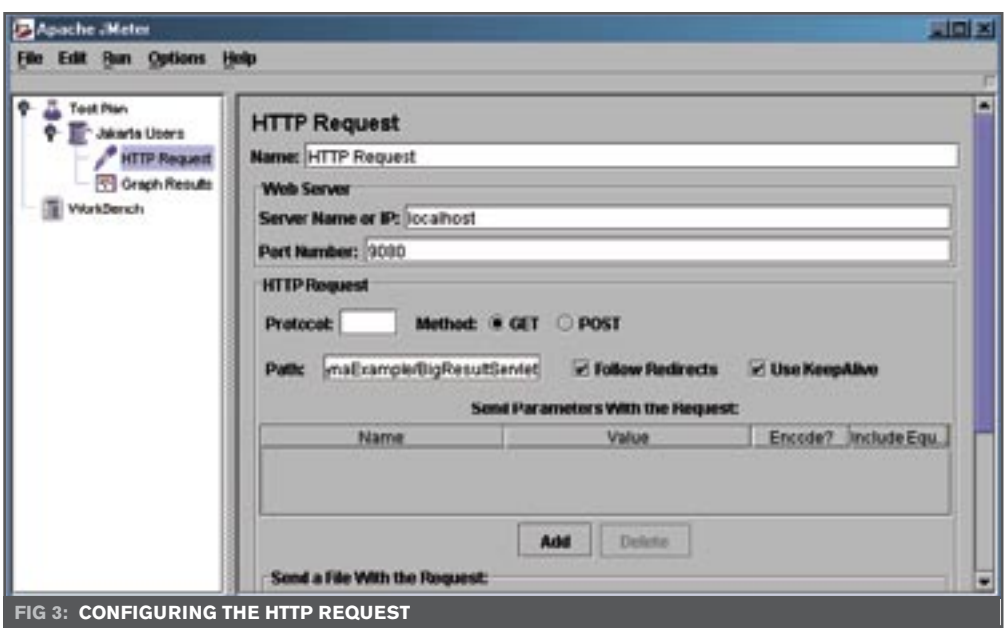

FIG 3: CONFIGURING THE HTTP REQUEST

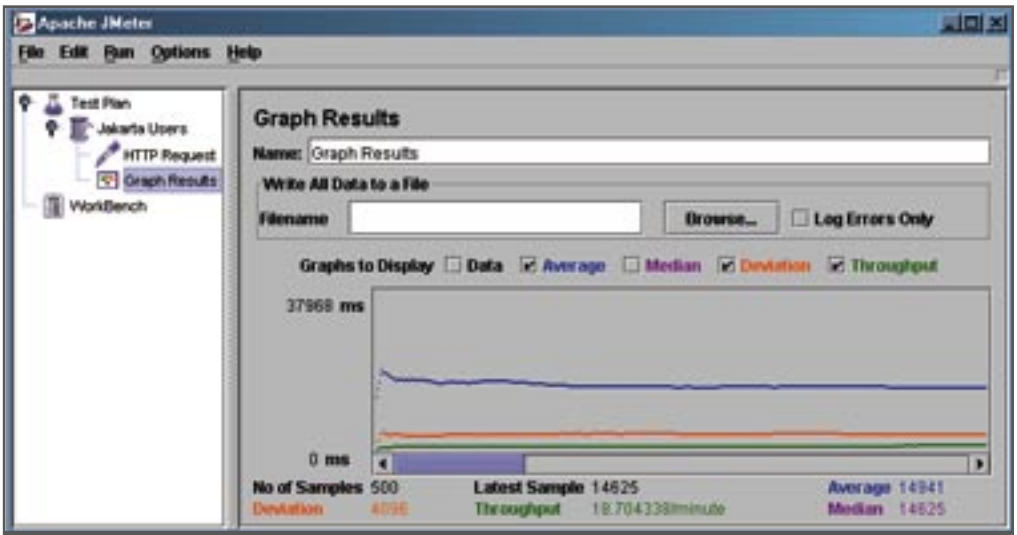

FIG 4: **RESULTS OF THE BASELINE TEST**

hybrid mode called shared-push-pull. These options are described in detail in the WebSphere Application Server InfoCenter.

- ***name:*** This is the URI of the servlet to cache; it can be relative to the Web application context root, as is ours.
- ***cache-id:*** This section describes how WebSphere should build the unique identifiers' cache entries. The components of the cache ID serve as rules that the cache manager uses to build a key value. For servlets, there are nine types of components that you can use to build a cache ID construction rule, giving you rich control over what content in your application is cached. In our example, we used components consisting of any parameters passed in on the servlet request, the URI path, plus a header field called "host" in the request. At least one of these components must be non-empty, otherwise the dynamic cache service will simply not cache the data for that request.
- ***timeout:*** By specifying a time-to-live value of 0, we specify that we want our cache entry to live indefinitely. Specifying a positive number would define the number of seconds we want the cache entry to exist in our cache.

We have only scratched the surface of the configuration parameters you can set for a cacheable object. For example, you can opt to persist cache entries to disk in case of overflows or the housing server being stopped. Refer to the WebSphere Information Center (www-306.ibm.com/software/webservers/appserv/infocenter.html) to find out more information about which parameters can be used to achieve the desired caching effect you'd like to associate with a specific cacheable object. These design considerations are of course specific to your enterprise application's business requirements.

The WebSphere Application Server continually monitors the cachspec.xml file for changes, so the changes will take place immediately when you save the file. You therefore don't need to restart your enterprise application or application server when you want to make adjustments to your caching strategy.

The Performance Gain with the Dynamic Caching Service Enabled

Once we saved our cachespec.xml file in the WEB-INF directory, WebSphere Application Server noticed the new file, and wrote the following line to System.out:

```
[2/2/04 17:09:50:812 CST] 435b6a5a ConfigManager I
DYNA0047I: Successfully loaded cache configuration file k:
\wsad\workspace5\Dynamic Cache Example\Web Content\
WEB-INF/cachespec.xml.
```

Then we reran our JMeter test. Figure 5 shows the results.

We started with a throughput of 18.7 transactions per minute, and with cachspec.xml in place we were able to achieve quite a bit more: 568 transactions per minute. The average time to serve a request went from 14.9 seconds down to a speedy 0.416 seconds. The dramatic improvement demonstrates the power of using WebSphere's dynamic caching service.

Our test servlet was deliberately constructed to demonstrate a dramatic effect, but most Web applications that do any processing or back-end calls will benefit substantially from

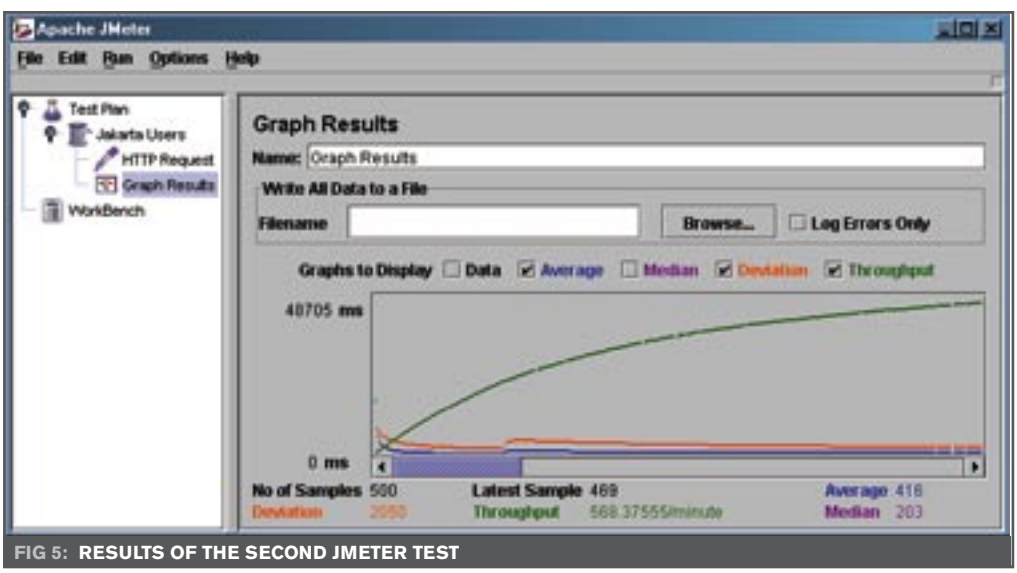

FIG 5: **RESULTS OF THE SECOND JMETER TEST**

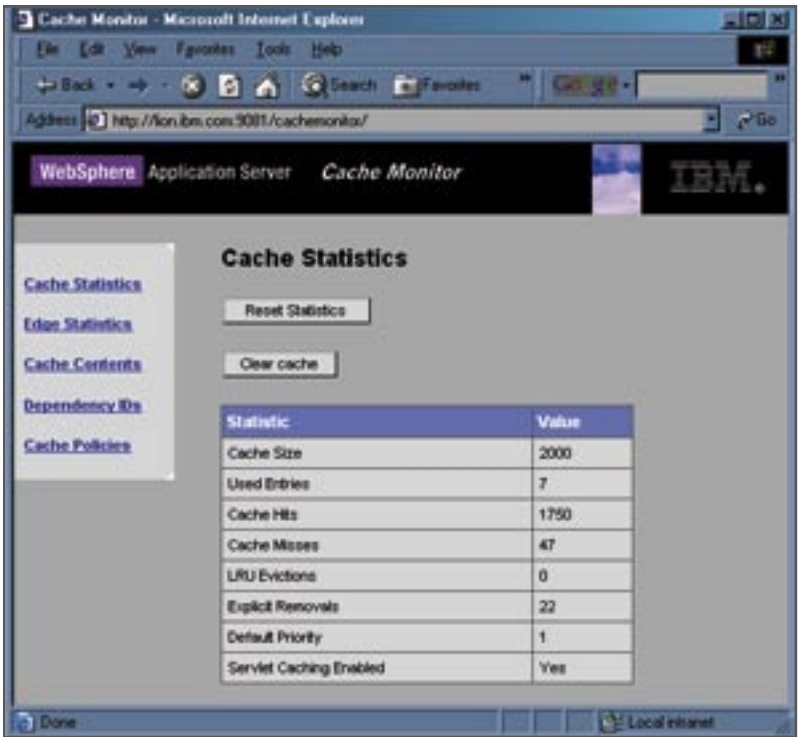

FIG 6: **THE CACHE MONITOR APPLICATION**

using the dynamic cache service. It's important to point out that we didn't need to make any changes to our servlet code. We simply created a cachespec.xml file in our WEB-INF directory and witnessed an immediate performance improvement.

WebSphere Application Server 5.0 ships with a Cache Monitor application (see Figure 6) that allows you to use a Web browser to view reports on cache statistics, and also allows you do things like view the cache IDs of objects in the cache.

The Cache Monitor is located in the <WAS_ROOT>/installableApps directory (CacheMonitor.ear) and is easy to install and use. After you have installed the Cache Monitor enterprise application, open the URL: http://<your_server>:<your_port>/cachemonitor. You can use the Reset Statistics and Clear Cache buttons to clean out the cache and get ready for another load test after you tweak the caching policy settings in your cachespec.xml file. This saves time, since you won't have to restart the application server between load tests.

Conclusion

Java servlets and JavaServer Pages are typically used to present dynamic content to a Web site visitor. WebSphere Application Server's dynamic cache service can be used to serve requests for content from an in-memory cache. Servlet output is stored in memory, expediting delivery to the end user.

In this article, you used the dynamic cache service offerings on a Web application and witnessed the added performance benefits using Apache JMeter. The process of enabling WebSphere Application Server's dynamic cache service and modifying your Web application to use dynamic caching is rather simple. Leveraging these highly configurable caching options should be a major consideration for organizations trying to squeeze the most performance out of their application servers. You can learn more about the dynamic cache service in the WebSphere Information Center.

"To see the effect of using dynamic caching in our application, we'll juxtapose the performance of our application with and without caching enabled"

Resource:

- *Using the dynamic cache service to improve performance:* http://publib.boulder.ibm.com/infocenter/wasinfo/index.jsp?topic=/com.ibm.wasee.doc/info/ee/ae/tprf_dynamiccache.html

Acknowledgment

The authors would like to thank Lisa Tomita for her review of this article.

Brad Bouldin works in the Portal Practice in IBM Software Services for WebSphere. bbouldin@us.ibm.com

Kulvir Singh Bhogal works as an IBM Software Services for WebSphere consultant, devising and implementing WebSphere-centric solutions at customer sites across the nation. He has over 50 patents pending in a myriad of technology areas. kbhogal@us.ibm.com

LISTING 1

```
package com.ibm.example;

import java.io.*;
import java.net.URL;
import java.util.Date;
import javax.servlet.*;
import javax.servlet.http.*;

public class BigResultServlet extends HttpServlet {

  public void doGet(HttpServletRequest req,
  HttpServletResponse resp)
        throws ServletException, IOException {
        System.out.println("BigResultServlet.doGet(
        ) starting");
        long start = System.currentTimeMillis();
        resp.setContentType("text/html");

        // open up a URL connection to our book
        URL bigData = new URL("http://
        www.gutenberg.net/etext96/olivr10.txt");
        sun.net.www.content.text.PlainTextInputStre
       am inputStream =
                (sun.net.www.content.text.PlainText
        InputStream) bigData.getContent();
        BufferedReader dataIn = new
        BufferedReader(new InputStreamReader(input
       Stream));
        ServletOutputStream output =
        resp.getOutputStream();

        output.println("<!DOCTYPE HTML PUBLIC \"-//
        W3C//DTD HTML 4.01 Transitional//EN\">");
        output.println("<HTML>");
        output.println("<HEAD>");
        output.println("<title>Big Result Servlet
        - Dynamic Cache example</title>");
        output.println(
                "<meta http-equiv=\"Content-Type\"
        content=\"text/html; charset=iso-8859-1\"
        />");
        output.println("<BODY>");

        output.println("<PRE><I>");

        String line = dataIn.readLine();
        while (line != null) {
                // if we reach the end of the
                // preamble, stop italicizing
                // and switch the font
                if (line.endsWith("*END*")) {
                        output.println(line);
                        output.println("</I><font
face=\"Times New Roman\">");
                } else {
                        output.println(line);
                }
                line = dataIn.readLine();
        }

        output.println("</PRE></FONT>");
        output.println("</BODY></HTML>");

        long end = System.currentTimeMillis();
        System.out.println("BigResultServlet.doGet(
) duration in millisec:" + (end - start));
  }
}
```

LISTING 2

```
<?xml version="1.0" ?>
<!DOCTYPE cache SYSTEM "cachespec.dtd">

<cache>
  <cache-entry>
      <class>servlet</class>
    <sharing-policy>shared-pull</sharing-policy>
      <name>/BigResultServlet</name>
      <cache-id>
         <component id="*" type="parameter">
            <required>false</required>
         </component>
         <component id="" type="pathinfo">
            <required>false</required>
         </component>
         <component id="host" type="header">
            <required>false</required>
         </component>
         <timeout>0</timeout>
      </cache-id>
  </cache-entry>
</cache>
```

Versata's Rules

for

Simplifying WebSphere Studio Development:

RULE THREE

DEVELOPERS LIKE NEW SOFTWARE, ESPECIALLY **FREE** SOFTWARE.

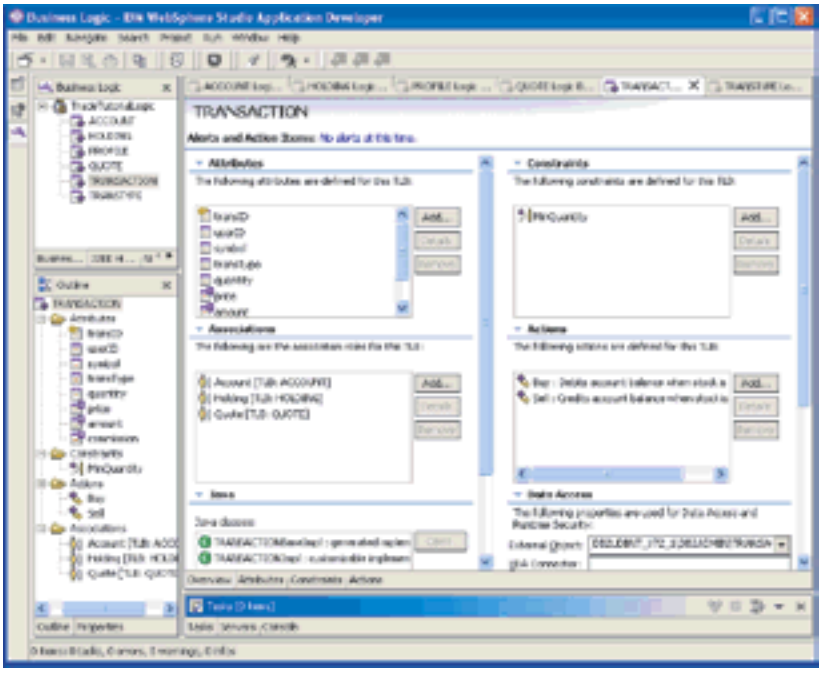

J2EE can be a bear without Versata business logic software. As a plug-in for IBM's WebSphere Studio, Versata helps you develop and manage server-side business logic components, driving data for JSPs, HTML, Java, and Web services - making life a whole lot easier. Rather than program hundreds of lines of code, you simply specify your business logic as business rules: Versata does the rest.

versata™

To download a **free** copy of the Versata Business Logic Designer for WebSphere Studio, visit **www.versata.com/businesslogicdesigner**

Secure Messaging in a Dangerous World

Creating a trustworthy mechanism for information exchange

BY JEFF **DEMENT**

With the rise of asynchronous messaging comes the need for securing message flows in WebSphere MQ. This article discusses the architecture of WebSphere MQ, along with some options available for securing an asynchronous communications infrastructure.

Today's business environment is becoming more and more complex, and the demands on the IT department to support increasing numbers of enterprise initiatives are growing in lockstep. One area drawing significant attention is communications, both between applications within the enterprise and between applications that span companies. While communications over the Internet have traditionally been characterized in terms of HTTP (browser/Web server, and more recently, Web applications running over SOAP), other communication styles are often more efficient and/or more appropriate.

Asynchronous messaging is a popular and effective communications paradigm for integrating applications within the enterprise as well as facilitating electronic business relationships that span enterprise boundaries. Nevertheless, as with all forms of communications (see sidebar on XML), security is the critical requirement – and the level of trust that can be established is often the major determinant of the business value the integration returns to the enterprise.

In this article I will describe the integration of X.509 digital certificates and messaging to create a trustworthy mechanism for the exchange of sensitive business information. In particular, I will describe in some detail the architecture of IBM's WebSphere MQ (formerly known as MQSeries), with an emphasis on the options for securing message flows using capabilities enabled by aspects of VeriSign Intelligence and Control Services – public key cryptography, digital certificates, and public key infrastructure (PKI).

Introduction to IBM WebSphere MQ

I will describe a secure asynchronous communications infrastructure. WebSphere MQ commands between 70–80 percent of this market and will be the sole focus of this article. This section provides a gentle WebSphere MQ introduction; a more detailed description comes later.

With deep apologies to IBM development, WebSphere MQ can be likened to an e-mail system on steroids. Although this analogy is a gross simplification, an e-mail system has a message sender and receiver (neither must be online or connected for the other to do useful work), both sender and receiver have a mailbox, to hold messages pending transmission or processing respectively, and sender and receiver are connected through a communications medium (most likely TCP).

Figure 2 illustrates a much-simplified WebSphere MQ architecture (analogs to an e-mail system are obvious; queues equate to mailboxes and the queue manager/MCA [message channel agent] combination functions much like the sendmail daemon):

- ***Queue manager:*** Controls the WebSphere MQ environment on a particular machine, facilitating message transfer from the business application through a queue, onto a communications channel vis-a-vis the message channel agent.
- ***Queue:*** A (temporary or permanent) holding area for messages (i.e., before they are put onto the communications medium at the sending machine, or consumed by the application on the receiving machine).
- ***Message channel agent:*** Deals with the complexity of the communication protocol and manages the guaranteed delivery of messages from the sending machine to the receiving machine.
- ***Channel:*** The pipe that connects two communicating WebSphere MQ machines – typically TCP, but other protocols (such as SNA) are supported.

Message queuing is a method of application-to-application communication. Applications communicate by writing/reading application-specific data (messages) to/from queues, without having a dedicated synchronous link between them. Messaging means that programs communicate with each other by sending discrete amounts of data in messages to some intermediary, and not by calling each other directly (unlike the RPC paradigm – remote procedure calls). Queuing implies

that applications communicate through queues, which removes the requirement for both sending and receiving applications to execute concurrently. In other words the sending and receiving of messages is asynchronous; there is no time dependency between sending and receiving except that which may be imposed by the applications themselves.

WebSphere MQ is IBM's implementation of messaging. WebSphere MQ–enabled applications use a consistent Application Program Interface (API) across all platforms. The application communicates directly with its local queue manager by using the Message Queuing Interface (MQI). The MQI is a small API with only 13 calls that request services from the queue manager.

The typical programmatic flow of WebSphere MQ is shown in Figure 3. The first step is for the application to connect to the queue manager. It does this through the MQConnect call. The next step is to open a queue for output using the MQOPEN call. The application then puts its data on the queue using the MQPUT call. To receive data the application will also call the MQOPEN call to open an input queue. The application will receive data from the queue using the MQGET call.

Also shown are the message channel agent, channel exits, and Object Authority Manager (OAM). The MCA is the WebSphere MQ program that moves the messages from the local transmission queue on the sending machine to the target queue manager on the receiving machine using existing transport services. The most commonly used transport services are TCP/IP and SNA. The channel exits are user-written programs that can be executed from one of a defined number of points in the program execution flow during channel operation. The OAM is the default authorization service (which is operating system specific) for command and object security. These three components are important for existing security solutions for WebSphere MQ.

EXISTING SECURITY CAPABILITIES WITHIN WEBSPHERE MQ

Existing security capabilities within WebSphere MQ are in three components: the OAM, channel exits, and (available with the current release, v5.3) SSL channels.

1. ***The Object Authority Manager (OAM):*** WebSphere MQ uses the OAM to provide access control functionality (i.e., authorization services). The OAM maintains an access control list for each resource that it controls. The OAM is passed a principal, a resource name, and an access type request. It will either grant or reject access based on its internal access control list.

2. ***Channel exits:*** WebSphere MQ channels provide three exit points where security functionality may be added (in fact, these exits can be used to implement any desired functionality, for example message compression):
   –*Security exit:* Used for authentication of the partner MCA. The exit is invoked at MCA initiation and termination, and at channel startup after initial negotiation, but before any user data is exchanged. These channel exit routines operate in pairs.
   –*Message exit:* Used for link encryption (such as SSL), validation of user ID, and auditing. A message exit is invoked once per message on each side of the channel, that is, after the message is taken off the transmission queue and before the message is put on the target queue.
   –*Send and receive exit:* Used for data encryption and compression. These exits are called once per message segment. The data available to the exit program is part of the message.

   Channel exits are commonly used to protect data. They can provide data encryption, data integrity, and message origin authentication. The data protection is on a machine-to-machine basis and operates at the transport (communications protocol) level, not the application layer.

3. ***SSL channels:*** New in WebSphere MQ v5.3 is a capability called SSL channels. As the name implies, this is functionality that implements standard SSL over TCP/IP. This feature provides channel authentication, data encryption, and data integrity for MQ messages on the wire (i.e., as they travel from MCA to MCA over a

XML, A MIXED BLESSING

It is important to understand the (perhaps not so) subtle impact of using XML to encode WebSphere MQ messages. Figure I shows a string of digits being snooped by an unauthorized individual eavesdropping on the channel, but without a detailed knowledge of the sending or receiving application and more specifically the data formats being exchanged, it is difficult to ascribe any meaning to the top flow.

On the other hand, the identical data flowing in the XML self-describing format makes the job of data interpretation much easier. XML places business context (i.e., meaning) around an otherwise incoherent sequence of numbers. XML facilitates many benefits to a corporation (such as application decoupling), nevertheless all other things being equal, XML may increase the security requirements for transmitting confidential data.

FIG 1: XML, A MIXED BLESSING

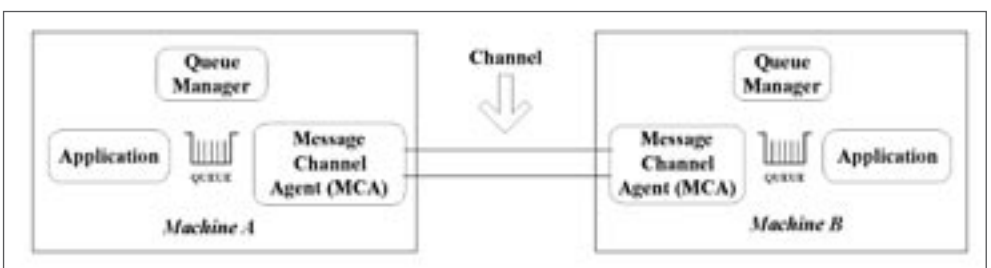

FIG 2: SIMPLIFIED MESSAGING ARCHITECTURE

TCP connection) but has no effect on messages in the sending or receiving queues.

Architectural Perspective on Enhancing WebSphere MQ Security

This section describes potential integration points for enhancing or extending WebSphere MQ security from an architectural perspective. Figure 4 highlights the specific points where security can be integrated.

MODIFY THE APPLICATION

Integration Point 1 positions security policy (implemented with programming logic) within the application, which is generally considered a bad idea from a security standpoint. On the other hand, it also means the application has complete control over which security services are provided and at what granularity.

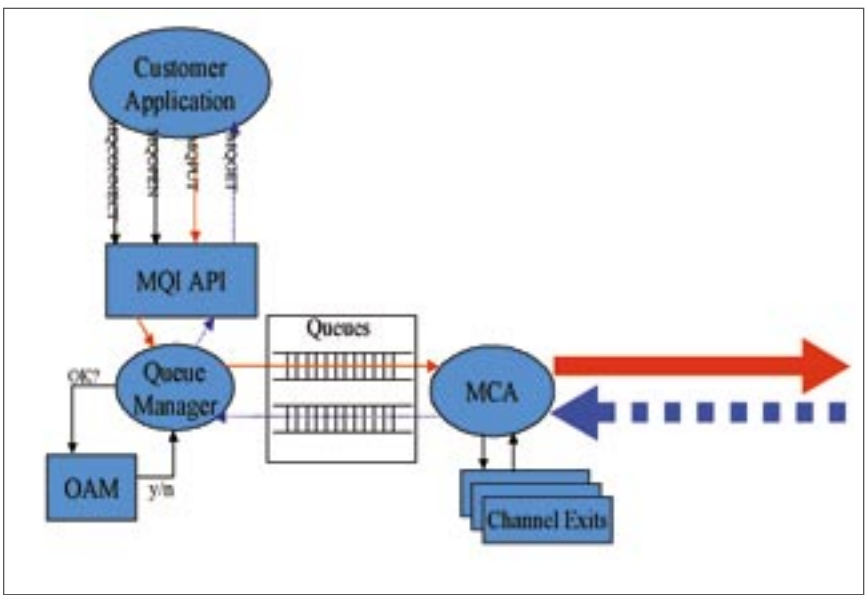

FIG 3: HIGH-LEVEL WEBSPHERE MQ ARCHITECTURE

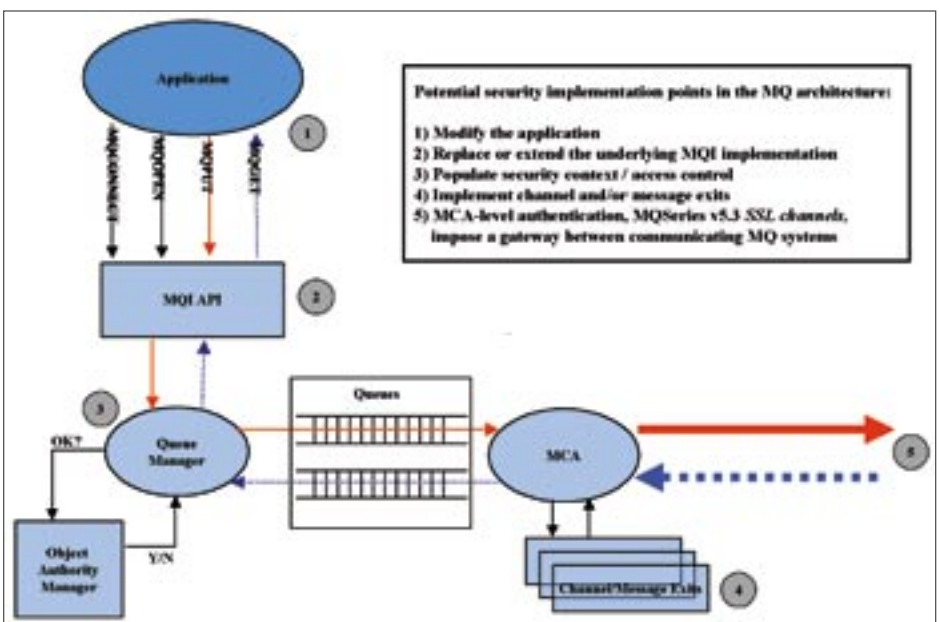

FIG 4: WEBSPHERE SECURITY INTEGRATION POINTS

REPLACE THE MQI LIBRARY OR MODIFY ITS BEHAVIOR

Integration Point 2 can be implemented with either of two fundamental approaches:

1. Replacing (or displacing) the standard WebSphere MQ library that provides the MQI interface with a new library that provides security services
2. Extending the MQI by taking advantage of WebSphere MQ API Exits, a general-purpose framework for modifying the behavior of any MQI call

The first approach would generally represent a vendor-provided solution, while the second could be a vendor-provided solution or a customer roll-your-own alternative.

OBJECT AUTHORITY MANAGER

Integration Point 3 is included in the interest of completeness and represents existing OAM functionality, i.e., authorization to WebSphere MQ resources. On the mainframe equivalent OAM functionality is implemented with RACF (Resource Access Control Facility).

CHANNEL EXITS

Integration Point 4 also represents an existing WebSphere MQ security capability, although out-of-the-box exits are simply exit points (architectural placeholders) that customers can use to implement security either by purchasing a channel exit–based security product or by writing one themselves.

CHANNEL OVERLAY

Integration Point 5 implies the imposition of a security service at the channel level, essentially tunneling WebSphere MQ messages inside a VPN-style service.

Practical Solutions to Enhance WebSphere MQ Security

This section maps various practical solutions into the architecture by describing some of the WebSphere MQ security products available today from IBM.

MS8I – WEBSPHERE MQ INTERNET PASS-THRU (MQIPT)

WebSphere MQ Internet Pass-Thru (MQIPT) is a WebSphere MQ base product extension (freely available SupportPac) that can be used to implement messaging solutions between remote sites across the Internet. It makes the passage of WebSphere MQ channel protocols into and out of a firewall simpler and more manageable by tunneling the protocols inside HTTP or SSL or by acting as a proxy.

If the SSL option is chosen, this product fits cleanly into Integration Point 5 in Figure 4, intercepting WebSphere MQ message flows and tunneling them through an SSL connection.

PREVIEWING MAY 11, 2004, AT NETWORLD+INTEROP, LAS VEGAS

INFORMATION STORAGE+SECURITY JOURNAL!

FOR MORE INFORMATION VISIT WWW.ISSJOURNAL.COM

SYS-CON MEDIA

From the World's Leading i-Technology Publisher

WEBSPHERE MQ V5.3 SSL CHANNELS

WebSphere MQ v5.3 has added new security functionality that provides (out-of-the-box) the ability to transmit messaging through a secure SSL connection. There is no intervening gateway; this is a new channel type built into the base product that provides for encryption of all messages flowing on this channel. This option also represents a solution that works at Integration Point 5.

WEBSPHERE MQ V5.3 API EXITS

WebSphere MQ API exits let you write code that changes the behavior of WebSphere MQ API calls such as MQPUT and MQGET. API exits are conceptually identical to the well-known WebSphere MQ channel exit capability discussed earlier, but operate at a much lower level in the WebSphere MQ code stack. The notion is to insert code immediately before or after any MQI call, and then the queue manager drives this exit code to the registered entry points. This technology works at Integration Point 2 in Figure 4.

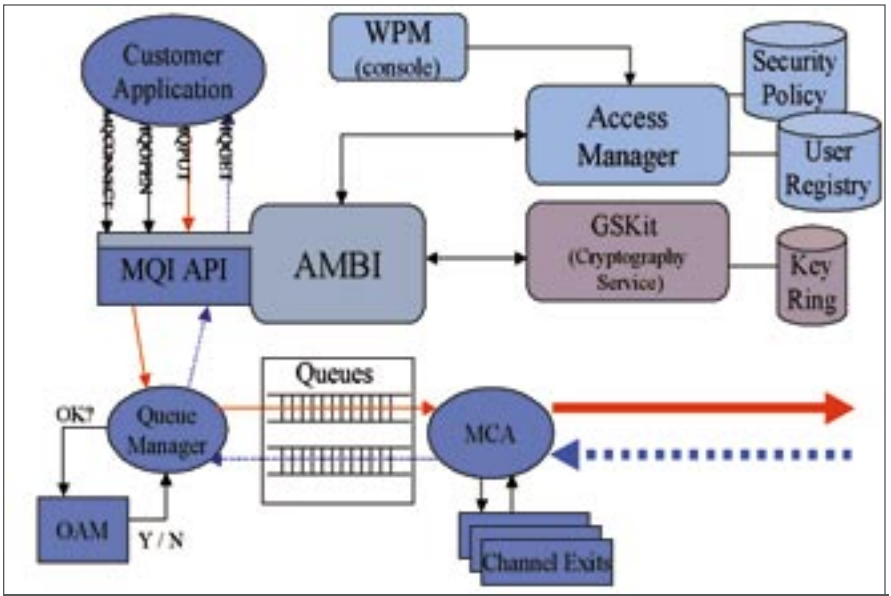

FIG 5: TAMBI ARCHITECTURE

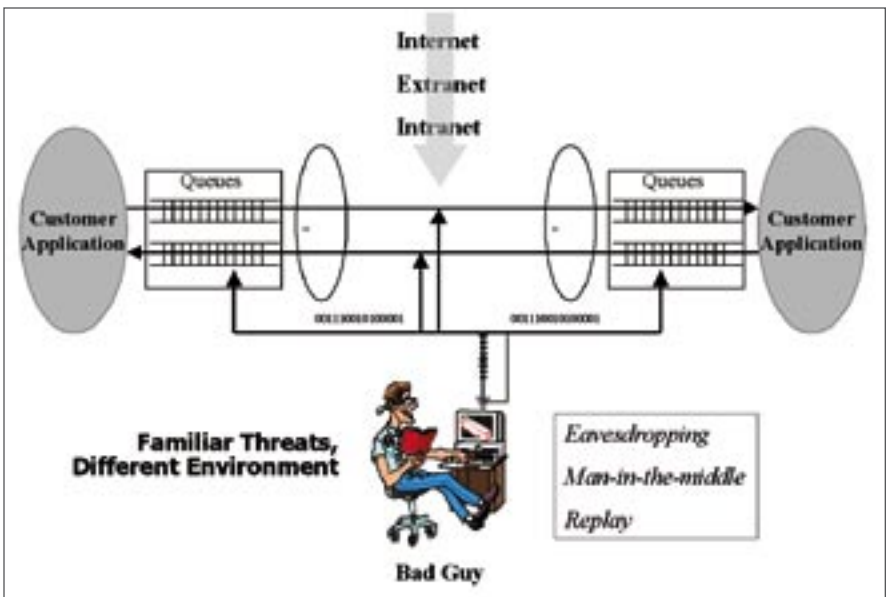

FIG 6: WEBSPHERE MQ PROBLEM SPACE

Access Manager for Business Integration

IBM Tivoli Access Manager for Business Integration (TAMBI) also fits into the WebSphere MQ architecture at Integration Point 2; this is clearly shown by its component architecture (see Figure 5).

TAMBI provides a set of multithreaded, shared libraries that execute in the process space of each WebSphere MQ application, and a per-machine daemon for managing functions common across all MQ-enabled applications on that machine. TAMBI libraries intercept WebSphere MQ API calls, enabling WebSphere MQ applications to be secured without modification. Calls made by the application to WebSphere MQ for services are captured by the interceptor, which performs the following four core functions:

1. An authorization check on the request for WebSphere MQ services (i.e., should this action be permitted).
2. Based on security policy, encrypt and/or digitally sign the message before placing the message on the queue on the sending machine.
3. Based on security policy, decrypt and/or validate the signature on the message before presenting the message to the receiving application.
4. Audit the message flows based on the authenticated application (or user) PKI identity.

The interceptor obtains policy information from a master policy server that is managed by an administration console called the Web Portal Manager (a WebSphere Application Server-hosted application).

TAMBI uses public key cryptography to protect WebSphere MQ message data from viewing by unauthorized parties (encryption) and from undetected tampering (digital signatures). This data protection is applied while messages are in a queue, as well as on the wire traveling from WMQ machine to WMQ machine In other words, true application-to-application-level security. It also provides an audit trail of (configurable) WMQ events. Finally, because of its integration with Tivoli Access Manager, Access Manager for Business Integration provides a centralized, enterprise-wide authorization capability on all WebSphere MQ resources (queue managers and queues).

WebSphere MQ Security Options Comparison

Table 1 gives a short comparison of the major advantages and disadvantages of implementing some/all security services at the five WebSphere MQ security integration points. This analysis looks at both the architecture as well as the existing implementations discussed above.

Case Studies

The remainder of this article uses the concepts discussed so far in the context of solving two real-world customer problems. More specifically, these case studies demonstrate the use of a software offering called "IBM-VeriSign Solution for Secure

e-business Integration" (affectionately known as Trusted MQ) to create secure messaging environments for two customers.

TRUSTED MQ

Trusted MQ addresses a very familiar problem space (although not necessarily familiar to the business messaging world that WebSphere MQ serves) by marrying a collection of technology and expertise in security from two premier companies in this space, IBM and VeriSign:

- WebSphere MQ SSL channels
- TAMBI
- VeriSign Managed PKI
- Consulting services

Trusted MQ uses IBM PKI-enabled products and VeriSign X.509 digital certificates to provide authentication (application-to-application or channel-to-channel), data confidentiality through strong encryption, data integrity, digital signatures, and auditing capabilities – that can selectively be brought to bear to solve customer security problems in their WebSphere MQ infrastructure.

CUSTOMER A

Customer A's environment was built around a very common WebSphere MQ hub-and-spoke architecture, although the complication here was that the customer operated in all 50 states in the U.S. and had two processing centers – one on the East Coast and one on the West Coast.

All applications homed in one of the hubs, and all applications were hosted in one of the two data centers.

Due to the distribution of the data, most WebSphere MQ traffic was confined to one (or the other) data center, but a small percentage – which still amounted to thousands of messages per day – required access to application(s) on the other hub. This traffic was initially sent in clear text that represented a clear exposure to the customer. In addition, because each hub represented a mission-critical processing point in the enterprise it was mandatory that only the New York hub was allowed to communicate with the LA hub, and vice versa. Rogue applications could not be allowed to connect to a hub, possibly impacting availability or worse.

WebSphere MQ SSL channels and VeriSign were chosen to solve this customer's security problems by providing the following features:

- SSL v3 using mutual authentication. One queue manager on each hub was provided a VeriSign X.509 digital certificate; its partner on the other hub validates the certificate (including certificate revocation list checking) and checks for an authorized distinguished name before allowing the channel to start.
- SSL provides the strong encryption needed to assure the necessary data confidentiality for the customer's sensitive data as it traverses the Internet.

CUSTOMER B

Customer B had a much more modest WebSphere configuration, only three machines arranged in a bus architecture. However, all three applications exchanged messages representing high dollar value business trans actions. Volume was low, but the data was extremely sensitive and needed to be protected at an application-to-application level.

In other words the message needed to be protected from unauthorized viewing or modification from the moment it left the sending application until the message was processed by the receiving application, including all time spent at rest in intervening queues.

Company audit requirements specified that every message needed to be audited (who sent it and when,

INTEGRATION POINT	ADVANTAGES	DISADVANTAGES
1 (Modify the application)	Flexibility; the application has complete control over the security functions provided and their granularity Messages can be protected on the queues as well as on the channel (end-to-end from a WebSphere MQ perspective)	Vendor lock-in because of the use of proprietary APIs Embeds security policy (security knowledge/rules) in the application Requires modification to the application at both ends to implement or modify security policy
2 (Replace or extend the underlying MQI implementation)	Application transparency; the application continues to code to the native WebSphere MQ API No security knowledge/rules in the application and security policy can be changed without impacting the application Messages can be protected on the queues as well as on the channel (end-to-end from a WebSphere MQ perspective)	Security is less granular than at the application level Introduces some unique service and configuration issues
3 (Object Authority Manager)	Out-of-the-box capability with WebSphere MQ Application transparency	Access control only Labor intensive No centralized administration or enterprise view of security
4 (Channel exits)	Application transparency Channel authentication is a one-time startup cost	Less flexible and granular Generally an all-or-nothing solution, not based on user, message, or queue Software distribution issues Messages in clear text while at rest in the sending/receiving queues
5 (Overlay security on the channel)	Application transparency Channel authentication is a one-time startup cost Performance will be better than performing security on a per message basis (1 and 2)	Less flexible and granular Generally an all-or-nothing solution; not based on user, message, or queue Messages in clear text while at rest in the sending/receiving queues

TABLE 1: ADVANTAGES AND DISADVANTAGES OF IMPLEMENTING SECURITY SERVICES AT THE FIVE WEBSPHERE MQ SECURITY INTEGRATION POINTS

who received it and when, how was it protected). The messages needed to be encrypted (eavesdropping by company employees was as much a concern as external hackers) and only authorized programs could be permitted to read/write messages from the queues.

IBM Tivoli Access Manager for Business Integration and VeriSign Managed PKI were chosen to meet these customer security requirements:

- TAMBI provides for digital signatures and encryption of each message (security policy is specified on a per/queue basis, in this case each application had only one queue).
- TAMBI audits each message based on the application's PKI identity (as defined in their X.509 certificate); these records contain a message identifier, sender or receiver identity, time stamp information, security policy applied, and other information.
- TAMBI provides a robust authorization service (access control to queues) guaranteeing that only the authorized application was allowed to PUT (write) or GET (read) messages from the respective queues.
- VeriSign's Managed PKI offering provided the software and services necessary for this customer to establish its own private certificate authority to issue certificates.

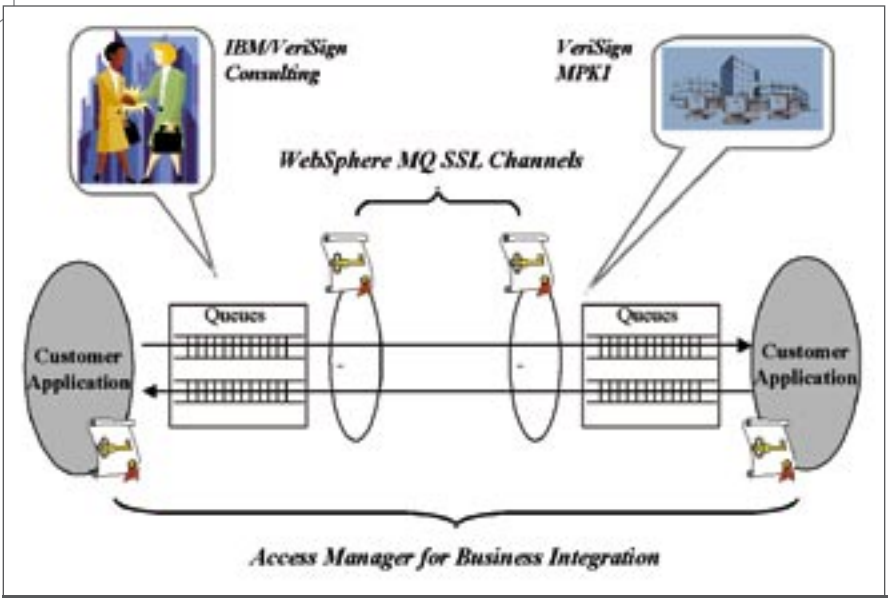

FIG 7: **WEBSPHERE MQ PROBLEM SPACE**

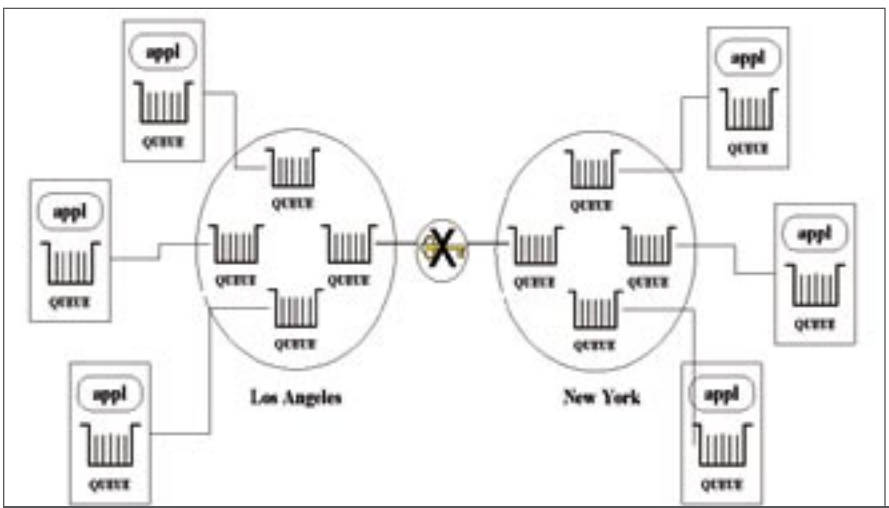

FIG 8: **CUSTOMER A ENVIRONMENT**

Summary

As with any technology that transmits potentially sensitive or high-value information over insecure networks (and they are all insecure at some level) multiple dangers exist:

- Unauthorized viewing of the information
- Unauthorized modification of the information
- Lack of auditability of the information flows
- Lack of assurance of who the communicating partner really is
- Unauthorized access to the communications medium

The architecture of WebSphere MQ was explored at some depth to illustrate the design points where security services could be integrated to address these dangers. Finally, real-world solutions were presented in the form of case studies using typical customer scenarios.

Resource

- *IBM WebSphere MQ family SupportPacs:* www- 306.ibm.com/software/integration/support/ supportpacs

Jeff DeMent has worked in the IT industry for 30 years, the past 5 at IBM providing security and directory solutions to customers. His recent focus has been on the Tivoli Access Manager family of security products, specializing in IBM Tivoli Access Manager for Business Integration, a secure messaging product. Recently Jeff has been involved in the VeriSign-Ready certification of TAMBI, as well as the new WebSphere MQ v5.3 SSL channels capability. Jeff skydives for fun, and has performed nearly 1600 jumps. jdement@us.ibm.com

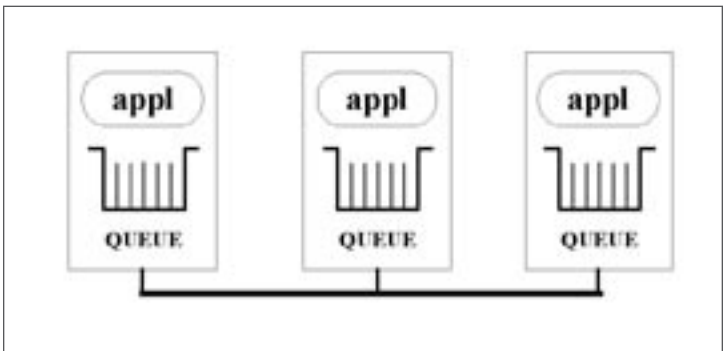

FIG 9: **CUSTOMER B ENVIRONMENT**

LOOK FOR YOUR FREE...

Coming this **SPRING!**

Using WebSphere Commerce to generate reports from archived user-traffic and sales data

Keep a Finger on the Pulse of Your Business

BY SAPNA MAHWAL & RAHUL KITCHLU

IBM WebSphere Commerce software provides powerful sell-side solutions to handle the challenges encountered in customer and trading partner environments. It can implement B2C, B2B, or private exchange business models using open, industry-accepted standards. With the introduction in WebSphere Commerce of hosted business models that handle extremely high volumes of traffic, it becomes especially important to archive sales and user-traffic data.

This article briefly describes the data back end of a WebSphere Commerce application, highlights a typical solution to archive useful information from this data, and displays that information using WebSphere Commerce.

Environment Details

A typical environment for such an application is a three-tier setup that contains an HTTP server at the front end to channel incoming requests, a middle tier powered by WebSphere Commerce, and a relational database management system, which hosts the back-end data and maintains persistence. Our solution uses IBM DB2 Universal Database.

IBM WebSphere Commerce delivers a market-leading B2B commerce solution to optimize customer and channel interactions, generate increased revenue, reduce costs and cycle times, and improve customer loyalty.

DB2 Universal Database is IBM's relational database management system for AIX, Linux, HP-UX, Sun, and Windows. DB2 database software marks the next stage in the evolution of the relational database: it is the industry's first multimedia, Web-ready relational database management system delivering leading capabilities in reliability, performance, and scalability.

IBM HTTP Server can be the foundation of any e-business application. IBM e-business software, such as the WebSphere family of products, is designed to operate with many popular Web servers.

The sample scripts were tested on WebSphere Commerce on the 2000 platform with DB2.

Overview

High-volume Web sites powered by WebSphere Commerce extract incoming user requests and store this data in a relational back end. This data is periodically cleaned with the help of scripts such as dbclean, which are supplied with the product. Although needed to optimize performance, these cleanups may cause loss of data needed for future business intelligence. We advise running simple scripts on the database before running the dbclean utility to archive historical usage records. This is valuable information that can help sales and marketing departments evaluate trends and determine business direction.

This article describes a method for extracting and archiving such data and using it to generate a WebSphere Commerce report. It provides sample SQL scripts to create and populate the archive tables and a simple report to view these archives. Though the example in this article is specific to the hosted store business models available with the latest version of WebSphere Commerce, you can easily modify it to suit other scenarios.

Process – User Traffic Logging in DB2

Enabling User Traffic Logging in a hosted WebSphere Commerce application populates the USRTRAFFIC table in the back-end database by recording user activity at a site level. For more information about logging, see the online help for WebSphere Commerce at www.ibm.com/support/docview.wss?uid=swg27001853.

For a typical site, the USRTRAFFIC table grows fairly rapidly. It needs to be cleaned periodically to maintain site performance. The dbclean utility provided with WebSphere Commerce removes all entries from the USRTRAFFIC table by default.

Solution – Data Extraction Script

Before running the dbclean script on the database, use the SQL script included in the source code (avail-

able for download from www.sys-con.com/websphere/sourcec.cfm) to easily extract the information to be archived. To have a complete and accurate set of historic user-traffic and sales data, run this script at the beginning of each month, and then run dbclean to clean up the database and optimize its performance. The script extracts user-traffic and sales data for the previous month. It is a simple example that you can easily customize for specific business scenarios.

This script first defines a table in the database called ARCHIVE_WCS. This table contains rows and columns extracted from the database schema used to store user data.

Once you have run the scripts the data will be defined as shown in Figure 1. *Note:* You can make changes to the script provided to modify this table according to your requirements.

Next, this script runs some predefined SQL data-extraction queries. Listing 1 creates the archive table, and Listing 2 populates the archive table. The following SQL command reorganizes the ARCHIVE_WCS table to optimize performance:

```
reorg table ARCHIVE_WCS inplace
allow read access notruncate table
resume
```

The following SQL command extracts records from the archive table:

```
select * from ARCHIVE_WCS;
```

This SQL command drops the archive table:

```
drop table ARCHIVE_WCS;
```

Using the Sample Report

Now that the user data has been extracted to the newly created table, we will walk you through the creation of a sample report program you can write to graphically view this archived data. The sample report program is available for download as part of the source file. To enable it on your system, take the following steps.

Note: In the following section, <INSTALL_ DIR> refers to the directory in which WebSphere has been installed; <instance_name> refers to the name of your Commerce instance.

1. Place the report pages in the appropriate directories. XML files should be placed in <INSTALL_ DIR>\CommerceServer55\xml\tools\reporting. JSP files belong in the <install_dir>\AppServer\installedApps\studio\<instance_name>.ear\CommerceAccelerator.war\tools\reporting directory.

2. Register the report xml resources in the file resources.xml. Navigate to the <INSTALL_DIR>\CommerceServer55\xml\tools\reporting\ directory and modify resources.xml by pasting in the following lines:

```
<resourceXML name="UserTrafficPatt
ernsReport" file="reporting/UserTr
afficPatternsReport.xml" />
<resourceXML name="UserTrafficPatt
ernsReportDialog" file="reporting/
UserTrafficPatternsReportDialog.x
ml" />
<resourceXML name="UserTraff
icPatternsReportOutputDialog"
file="reporting/UserTrafficPattern
sReportOutputDialog.xml" />
```

3. Add the text strings used in your report to the appropriate .properties file. Modify Reports_en_US.properties in the <INSTALL_DIR>\AppServer\installedApps\studio\<instance_name>.ear\properties\com\ibm\commerce\tools\reporting\properties directory by pasting in the lines shown in Listing 3.

4. Register the report "view" commands to the database by running the SQL code shown in Listing 4.

5. Add access control. For each of the views registered into the back-end database in the previous step, insert one row into the ACACTION table. Check the current maximum acaction_id value by running the

Columns | Keys | Check Constraints | Table

Column name	Data type	Status	Length	Precisi
STOREENT_ID	INTEGER	Exists	-	-
YEAR	INTEGER	Exists	-	-
MONTH	INTEGER	Exists	-	-
TOTALVISITS	INTEGER	Exists	-	-
TOTASESSIONS	INTEGER	Exists	-	-
DISTINCTUSERS	INTEGER	Exists	-	-
TOTALORDERS	INTEGER	Exists	-	-
GROSSSALES	DECIMAL	Exists	10	10

FIG 1: A VIEW OF THE ARCHIVE_WCS TABLE

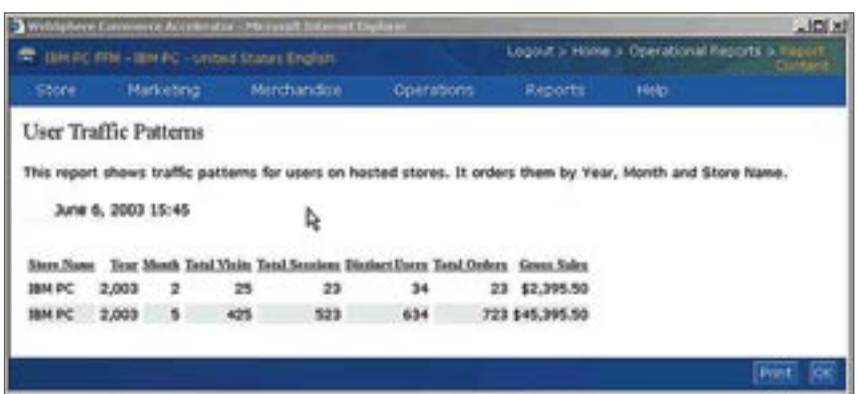

FIG 2: USER TRAFFIC PATTERNS REPORT

following SQL command in a DB2 command window:

```
db2 select max(acaction_id)
from ACACTION >
```

Our result was 11928. Now verify the current key table value for the table ACACTION by running the following SQL command:

```
db2 select counter from keys
where tablename='acaction' >
```

Our result was 11950. If you add enough rows to exceed the result obtained from the last step, you must update the keys table value for the table. For example, in our instance, if we had inserted enough rows into the ACACTION table to exceed the value of the counter in the table keys where tablename = acaction, we would have run the following SQL to update the keys table:

```
db2 select max(acaction_id)
from ACACTION
```

Our result was 11950. Increment the result obtained from the previous command to maintain uniqueness in the keys table. We changed our value to 11965.

```
db2 update keys set counter =
11965 where tablename = 'acac-
tion'
```

Run the SQL code shown in Listing 5 from a DB2 command window to add access control.

6. Add your report to the accelerator menu to access it:

```
Edit <INSTALL_DIR>\
CommerceServer55\xml\tools\
reporting\OperationalReportsCont
ext.xml.
```

Add the text shown in Listing 6 to the appropriate context tag, then add the text for the tag above to <INSTALL_DIR>\CommerceServer55\properties\runtime\com\ibm\commerce\tools\reporting\properties\OperationalReportsNLS_en_US.properties (see Listing 7).

7. Restart the WebSphere Commerce instance.
8. View the report (see Figure 2).

Summary

This article provided detailed instructions on how to consolidate and archive user traffic and sales data. Storing data in this form is more efficient and can improve the performance of your production site. This archived data was displayed in the GUI using the reporting framework provided with WebSphere Commerce. You can easily modify and extend this simple example to make it more robust.

Sapna Mahwal is a software specialist with IBM Software Services for WebSphere. Sapna has worked on the Commerce, Studio, and Services teams for the WebSphere family of products. smahwal@us.ibm.com

Rahul Kitchlu is a software engineer with the DB2 Technical Integration and Planning team at IBM's Toronto Software Lab, Canada. Rahul is an IBM Certified Solutions Expert, and his current focus is database technologies integration. rkitchlu@ca.ibm.com

LISTING 1

```
create table ARCHIVE_WCS
(STOREENT_ID int not null,
YEAR int,
MONTH int,
TOTALVISITS int,
TOTALSESSIONS int,
DISTINCTUSERS int,
TOTALORDERS int,
GROSSSALES decimal (10,2));
```

LISTING 2

```
insert into ARCHIVE_WCS
SELECT storeent.storeent_id as STOREENT_ID,
year(usrtraffic.shlustmp) as YEAR,
month(usrtraffic.shlustmp) as MONTH,
count(usrtraffic.usrtraffic_id) as TOTALVISITS,
count(distinct usrtraffic.sessionid) as TOTALSESSIONS,
count(distinct usrtraffic.users_id) as DISTINCTUSERS,
count( orders.orders_id) as TOTALORDERS,
coalesce( sum( totalproduct + totalshipping + totaltax + totaltax-
shipping + totaladjustment),0) as GROSSSALES
from storeent
left outer join usrtraffic on storeent.storeent_id =
usrtraffic.store_id
and month(usrtraffic.SHLUSTMP) >= month(current timestamp) - 1
and year(usrtraffic.SHLUSTMP) = year(current timestamp)
left outer join orders
on storeent.storeent_id = orders.storeent_id
and orders.status in ('M','A','B','C','F','G','R','S','D')
and month(orders.lastupdate) >= month(current timestamp) - 1
and year(orders.lastupdate) = year(current timestamp)
group by storeent.storeent_id,
year(usrtraffic.shlustmp),
month(usrtraffic.shlustmp);
```

LISTING 3

```
#
# User Traffic Patterns Report
#
UserTrafficPatternsReportWindowTitle=User Traffic Patterns Report
```

```
UserTrafficPatternsReportInputViewTitle=User
Traffic Patterns
UserTrafficPatternsReportDescription=This report
shows traffic patterns for users on hosted stores.
It orders them by Year, Month and Store Name.
UserTrafficPatternsReportCancelConfirmation=Select
OK to continue.

# Report output - input criteria
UserTrafficPatternsReportOutputViewTitle=User
Traffic Patterns Report
UserTrafficPatternsReportInputView=This report
shows traffic patterns for users on hosted stores.
It orders them by Year, Month and Store Name.
UserTrafficPatternsReportOutputViewTitle=User
Traffic Patterns

# Report output - table headings
UserTrafficPatternsStoreNameReport=Store Name
UserTrafficPatternsYearReport=Year
UserTrafficPatternsMonthReport=Month
UserTrafficPatternsTotalVisitsReport=Total Visits
UserTrafficPatternsTotalSessionsReport=Total
Sessions
UserTrafficPatternsDistinctUsersReport=Distinct
Users
UserTrafficPatternsTotalOrdersReport=Total Orders
UserTrafficPatternsGrossSalesReport=Gross Sales
```

LISTING 4

```
db2 insert into viewreg (viewname, devicefmt_id,
storeent_id, interfacename,
classname, properties, description, https, lastup-
date, internal)
values ('UserTrafficPatternsReportInputView', -1,
0,
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mand',
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mandImpl',
'docname=tools/reporting/UserTrafficPatternsReportI
nputView.jsp',
'', 1, CURRENT TIMESTAMP, 1)

db2 insert into viewreg (viewname, devicefmt_id,
storeent_id, interfacename,
classname, properties, description, https, lastup-
date, internal)
values ('UserTrafficPatternsReportOutputView', -1,
0,
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mand',
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mandImpl',
'docname=tools/reporting/UserTrafficPatternsReportO
utputView.jsp',
'', 1, CURRENT TIMESTAMP, 1)

db2 insert into viewreg (viewname, devicefmt_id,
storeent_id, interfacename,
classname, properties, description, https, lastup-
date, internal)
values ('UserTrafficPatternsReportDialogView', -1,
0,
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mand',
'com.ibm.commerce.tools.command.ToolsForwardViewCom
mandImpl',
'docname=tools/common/Dialog.jsp', '', 1, CURRENT
TIMESTAMP, 1)
```

LISTING 5

```
db2 insert into acaction (acaction_id, action) val-
ues (11929, 'UserTrafficPatternsReportDialogView');
db2 insert into acaction (acaction_id, action) val-
ues (11930, 'UserTrafficPatternsReportInputView');
db2 insert into acaction (acaction_id, action) val-
ues (11931, 'UserTrafficPatternsReportOutputView');
db2 insert into acactdesc (acaction_id, dis-
playname, description, language_id) val-
ues (11929, 'View to request the report for
UserTrafficPatterns', 'View to request the report
for UserTrafficPatterns', -1);
db2 insert into acactdesc (acaction_id, display-
name, description, language_id) values (11930,
'Specifies the criteria to create the report for
UserTrafficPatterns', 'Specifies the criteria to
create the report for UserTrafficPatterns', -1);
db2 insert into acactdesc (acaction_id, display-
name, description, language_id) values (11931,
'Displays the report for UserTrafficPatterns',
'Displays the report for UserTrafficPatterns', -1);
```

LISTING 6

```
<entry nameKey="UserTrafficPatternsReportName" des
criptionKey="UserTrafficPatternsReportDescription"
breadCrumbTrailTextKey="UserTrafficPatternsReport">
  <roles>
     <role>seller</role>
     <role>makMgr</role>
  </roles>
  <command name="UserTrafficPatternsReportDialog
  View">
     <parameter name="XMLFile" value="reporting.
     UserTrafficPatternsReportDialog" />
  </command>
</entry>
```

LISTING 7

```
#
# Bread Crumb Trail (BCT) Entries
UserTrafficPatternsReport=User Traffic Patterns

#
# Report Names
UserTrafficPatternsReportName=User Traffic Patterns

#
# Report Descriptions
UserTrafficPatternsReportDescription=This historic
report summarizes sales and user traffic data for
all stores on the site.
```

Integrating accounting applications using IBM WebSphere

The Adaptable Enterprise

BY GUIDO PUGLIELLI

Integrating heterogeneous subsystems and business units in a corporation is an increasing priority in a globalized economy in which the ability to adapt quickly in response to changes in market conditions determines the success or failure of an enterprise. This case study of an integration project for a large manufacturing/transport account details the use of components of the WebSphere framework.

From an IT perspective, there are many products, technologies, architectural design patterns, and methodologies available to accomplish the goal of adaptability. In this article I will present an application integration scenario, together with a solution based on technologies and products from the IBM WebSphere software platform.

Our customer (undisclosed) is a large manufacturing/transport account structured into several divisions – Holding, Human Resources, Finance, Services – and production departments. The business problem is to integrate these divisions, the services each one provides, and hence the different subsystems on which each division runs its own business.

As an instance of an integration scenario, let's consider the accounting business flow: suppliers, the goods/services acquisition approval cycle, orders, inventory, invoices, and accounting posting are the usual entities and flows involved in this scenario.

Several subsystems are involved in this kind of integration:

- SAP for materials management
- A third-party subsystem running General Ledger
- Several internally developed J2EE modules
- External business partners, such as banks

The requirements for this integration architecture were:

- From a business perspective:
  –To increase speed to comply with new government and industry regulations
  –To increase efficiency, reducing operational and resource costs
- From an IT perspective:
  –Asset reuse through a unified architecture
  –The enabling of nondisruptive additions of new applications to processes
  –The minimalizing of subsystem dependencies
- From a functional perspective:
  –Information integration through formatting rules and content-based routing
  –Application integration through standard protocols while trying to maximize information flow reuse
  –Business process integration, accomplished by combining several processes to produce a single flexible and extensible workflow

Application Integration Infrastructure

The solution discussed in this article was identified by taking into account the following general driving factors:

- Adoption of standard technologies, where applicable
- Adoption of market leader products in their relevant sectors
- Use of a single approach and of a single adapter technology for all application interactions with SAP, independent of the operating environment involved
- Extensive use of the following technologies:
  –Message-oriented middleware
  –Publish/subscribe
- The definition of application services to access EIS (Enterprise Information System) subsystems, incorporating parts of a service-oriented architecture in the solution, enabling product reuse in different scenarios; a loose connection between service providers and users; location transparency; independence from protocols, formats, and service life cycles
- Interaction based on events, incorporating elements of an event-driven architecture to provide excellent levels of scalability, robustness, and monitoring, along with zero coupling between the parties involved. (One party detects the change of state and publishes it, and all the services that have subscribed to it will receive the notification.)
- Definition of execution flows, obtained through human inter-

action and the invocation of the available application services; these flows, activated by events, model the company business processes and are easily modifiable in response to the evolution of company organization and strategies.

The IBM products involved in the resulting integration architecture (see Figure 1) and their features are described in the following sections.

WebSphere Application Server Enterprise Edition

WebSphere Application Server is IBM's J2EE-compliant engine; in its Enterprise Edition, infrastructure services and high-value application services have been added to the standard J2EE (EJB, JSP/servlet, JDBC, JCA, JMS) to make it a key platform for executing business processes. These include:

- The Process Choreographer's support for long-running workflows, including micro and macro flows, compensation, restart, and human interaction
- Support for Web services, protocols and standards, and the UDDI registry
- Support for messaging with point-to-point and publish/subscribe paradigms using JMS and MDB
- Strong characteristics of autonomy, performance, security, scalability, and system management

In this solution, WebSphere Application Server Enterprise Edition is the engine for executing the application workflows and for invoking the services on various EISs, e.g., SAP.

WebSphere MQ

WebSphere MQ (WMQ) is the de facto standard for messaging and queuing middleware. Its message delivery service, with its high quality of service, is available on almost every relevant hardware platform, and accessible from a broad range of programming languages, including Java, C, and COBOL; and operating environments, including J2EE and CICS.

In this architecture the application connectivity service is based on WMQ.

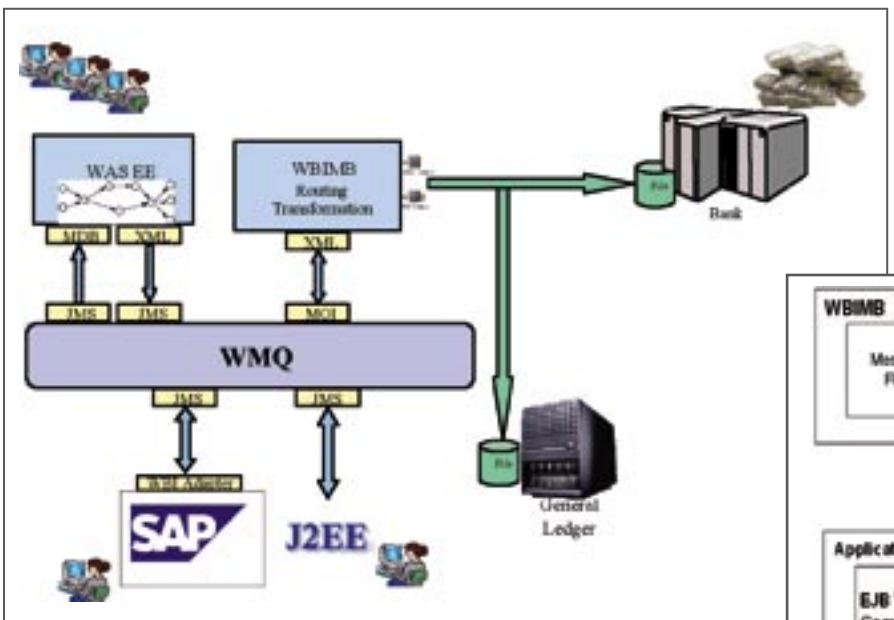

FIG 1: THE OVERALL INTEGRATION INFRASTRUCTURE: APPLICATION CONNECTIVITY SERVICES, PROCESS FLOWS, MESSAGE FORMATTING, FILE-BASED INTEGRATION

WebSphere Business Integration Message Broker

WebSphere Business Integration Message Broker (WBIMB) is a general and extensible integration broker that can manage data streams coming over WMQ or other protocols (e.g., HTTP), and that provides the following services:

- Routing based on content
- Format transformation
- Augmentation of content
- Topic- and content-based publish/subscribe
- Broadcast/multicast
- Support of custom operations through a complete plug-in architecture that can be implemented by the user

WBIMB provides native integration with WMQ. Both of these products allow you to implement high added-value interoperable application buses for all integration at the data level.

In this architecture, WBIMB is used to broker the exchange of data between the various subsystems involved in batch interactions or that request services such as routing and formatting.

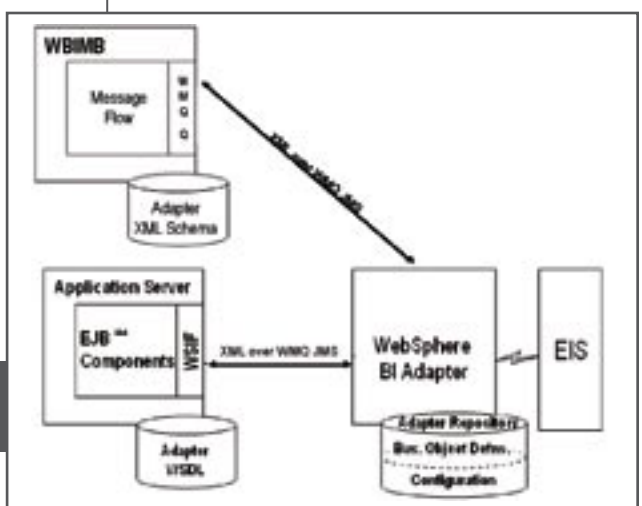

FIG 2: WBI ADAPTERS CAN BE ACCESSED FROM WBIMB AS WELL AS FROM THE APPLICATION SERVER

WebSphere Business Integration Adapter for mySAP.com

WebSphere Business Integration Adapters (WBI Adapters) are adapters for business applications and third-party middleware (EISs) based on CrossWorlds technology.

The internal architecture and the fundamental standard technologies are common to all adapters in the

WBI family. Each adapter contains a specific part relating to the EIS with which it must interface (e.g., mySAP.com); this part is typically based on an API and/or events provided by the EIS producer.

WBI Adapter services are available for the following operating environments:

- WAS
- WBIMB
- WBI InterChange Server

These adapters can take on a passive or active role:

- They can be activated by WAS or WBIMB to execute an operation on the EIS where they operate.
- They can be programmed to detect a modification of the state information managed by EIS, and as a consequence, generate an event.

In this architecture WBI Adapters are used to deal with mySAP.com from WBIMB as well as from WAS (see Figure 2).

Some Integration Scenarios – Process Flow

Example scenario: when an invoice arrives in paper form, an administration employee activates the workflow in order to manage it. The invoice data is entered via a browser, the relevant order is verified, receipt of goods or use of a service is verified, and regulation of payment is managed by a separate workflow. When feedback on payment is received, the invoice is archived and the payment is registered in the accounts subsystem.

Integration between procedures and applications is implemented through execution flows, which are created in a declarative rather than a programming way using the appropriate language (FDL [Flow Definition Language]) and development environment (the WebSphere Studio Integration Edition).

The flows are executed in WAS by the Process Choreographer component. Process flows can be activated by:

- Business events, generated by the SAP adapter when a modification to the state has been detected
- Scheduling, i.e., by defining rules for automatic flow activation in WAS
- Human intervention, e.g., when a new paper invoice is received from a supplier an employee starts a new flow instance to manage it

A business flow is typically composed of one or more tasks, each involving a human interaction or a process activity:

- Human interactions are accomplished through a customizable Web GUI (namely the process Web client) included in WAS; users authorized to operate on these flows are classified into roles and operate on tasks assigned to them executing basic "task claim" and "task completion" operations.
- Process activities are programmed using the services and standards provided by the WAS programming model, such as EJBs and Web services.

Application Connectivity with SAP

Example scenario: verification of the order associated with the invoice received – and verification of the existence of goods in inventory are typical services implemented by SAP, and these can be exposed and used in process flows.

SAP is accessed through services: the services provided by the WBI Adapter are described in WSDL and invoked by WAS via XML/JMS through the WSIF (Web Services Invocation Framework). These services can be used in a process flow executed in WAS.

The services approach is complemented by the adoption of a transport mechanism (WMQ) with assured and guaranteed delivery attributes:

- For insert/update/delete interactions (with side effects or business logic invocation) the use of persistent messages guarantees that the operation will be completed even in the case of a node failure, increasing the robustness of the internal infrastructure.
- For retrieval interactions or operations, in which double execution after the restoring of a failed node does not cause any undesirable effects, using nonpersistent messages improves performance without compromising robustness.

Event-Driven Interactions

Example scenarios: when a new item arrives it is registered in the inventory; this operation is published as an event, which causes a workflow to start in WAS. In another example, a suspended invoice process can finally be completed.

The events are detected and/or generated by WBI Adapters in the case of SAP, for

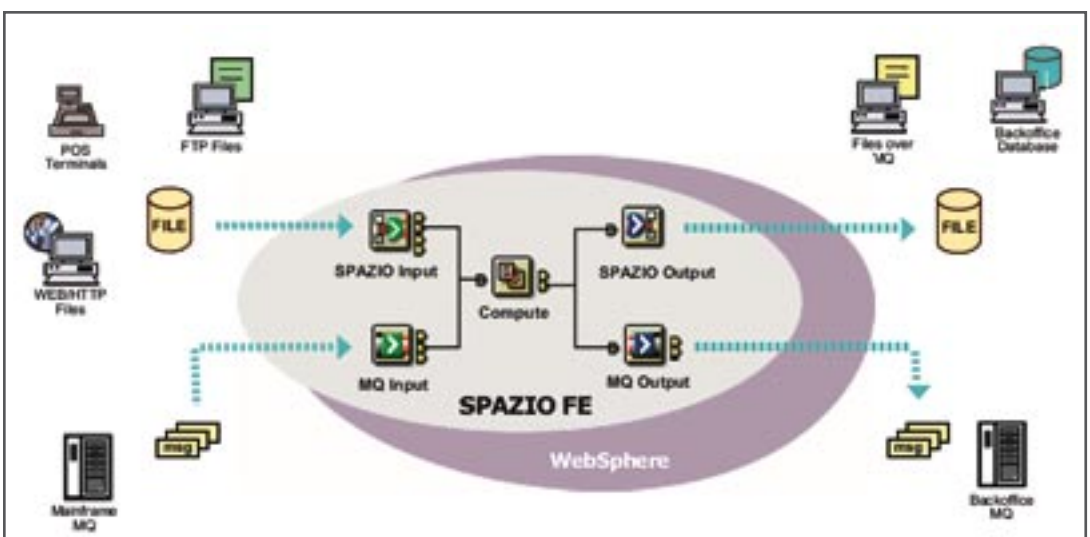

FIG 3: SOLUTION IMPLEMENTING PRIMEUR PRODUCTS

example, or by third-party applications that are involved in the architecture.

The events are transported by the application connectivity subsystem (WMQ) as messages. With its characteristics of secure and guaranteed delivery and the use of persistent messages, WMQ allows the activity to be completed even in cases of node failure.

The events are addressed to WBIMB, which transmits them, for example, via JMS/MDB to all the subsystems that have subscribed to receive them, carrying out any required processing and transformations to the transported data.

Data Manipulation and Batch Integration

In each phase of its life cycle, data is treated as a business object and not as an opaque byte stream; this is made possible through the virtualization of the base information provided by the WBI Adapters and through WBIMB parsing.

Data transmitted with standard protocols and formats such as JMS and XML can be manipulated in various ways; it is possible to adapt its format to the needs of the receiver or combine it with data retrieved from EISs (e.g., SAP or J2EE modules), and send it to file-based applications.

Example scenario: toward the end of each month it is necessary to activate the payment process of the invoices/bills that are about to expire. At a scheduled time a process flow starts in WAS: a summary of all payments to be made is displayed via a browser; when the user confirms the operation, the data regarding the payment to be carried out is sent to WBIMB, processed, and put into files that are sent to the different banks as appropriate.

This solution can be implemented with the help of products that build on WBIMB's data manipulation capabilities, in this case extended with file tranfer and batch features through Primeur's Spazio File Extender for WBIMB (see Figure 3).

Conclusion

This project, part of which is already in production and part of which is in the process of being completed, has provided several advantages, including increased efficiency, reduced operating costs, and the ability to reuse services and data. In addition, the solution is scalable and has been quick to meet business requirements.

Guido Puglielli is an IT architect for Primeur (www.primeur.com) with 10 years of experience in the IT industry. He is a WebSphere Certified specialist, is IBM Certified for e-business, and usually supervises development teams involved in integration projects using IBM WebSphere products and technologies. g.puglielli@primeur.com

Make your applications multilingual with WebSphere Studio

It's a Small World

BY KULVIR SINGH **BHOGAL**

Reality check: your code likely exists beyond the confines of your cubicle. In today's world – made smaller via the marvels of modern society – it is likely that your code will transcend borders or be seen by individuals who might not necessarily understand English. Not catering to these individuals can close the doors on a huge untapped market.

Fortunately, WebSphere Studio Application Developer v5 makes it easier to cater to crowds that speak different languages. As you'll see in this article, the concept of internationalization support is built directly into WebSphere Studio.

Teach Your Application to Speak Different Languages in Minutes

To showcase the internationalization support in WebSphere Studio, we'll play around with the class shown in Listing 1. Start by creating a Java project in WebSphere Studio to house the class. When you are done, right-click in the editor view and choose the Source>Externalize Strings... option.

At this point, you will see an Externalize Strings window. Here you will be presented with a list of the Strings within your class. WebSphere Studio automatically generates a key value for each String. Click on one of the keys to see how it correlates with the Strings in your code (see Figure 1). You'll see in just a bit how this key value is used. You can choose not to have certain Strings translated by clicking the "Never Translate" button next to a given String. For demonstration purposes we will allow each of the Strings to be translated. Click Next.

In the next screen, you will be allowed to set a name for a properties file in which you will store the value a given key will translate to. We'll keep the default name, which is "test.properties". In this screen you can also specify the class name of the resource bundle access class that will be used to facilitate extraction of Strings from the properties file. Keep the defaults provided and click Next.

In the next screen, you will be informed that the file test.properties does not exist and will be created. Click Next.

In the final screen (see Figure 2), WebSphere Studio allows you to review the proposed changes before committing to refactoring. Note that

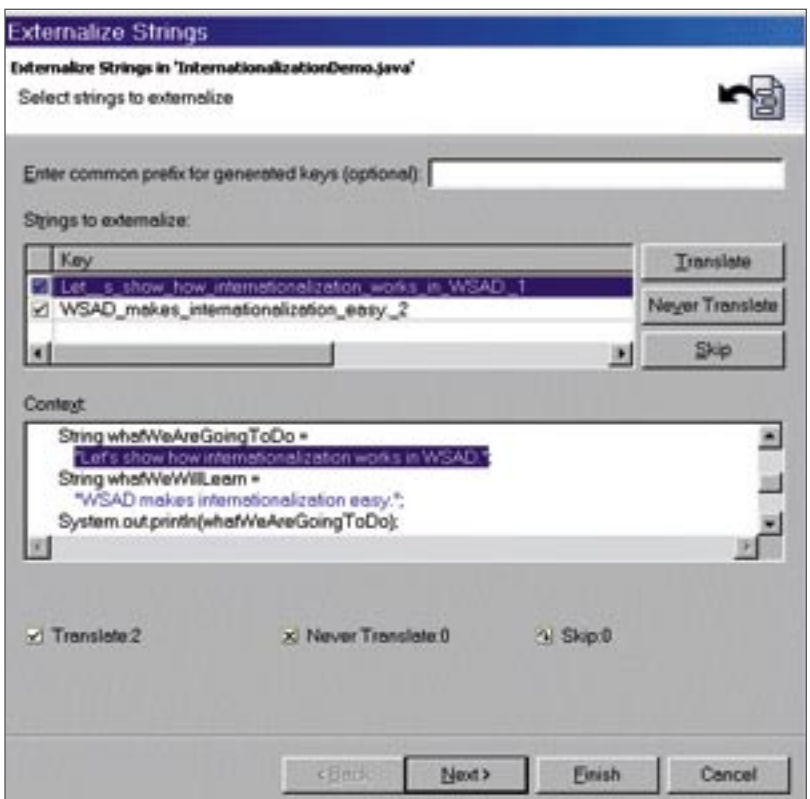

FIG 1: KEYS ARE AUTOMATICALLY GENERATED BY WEBSPHERE STUDIO

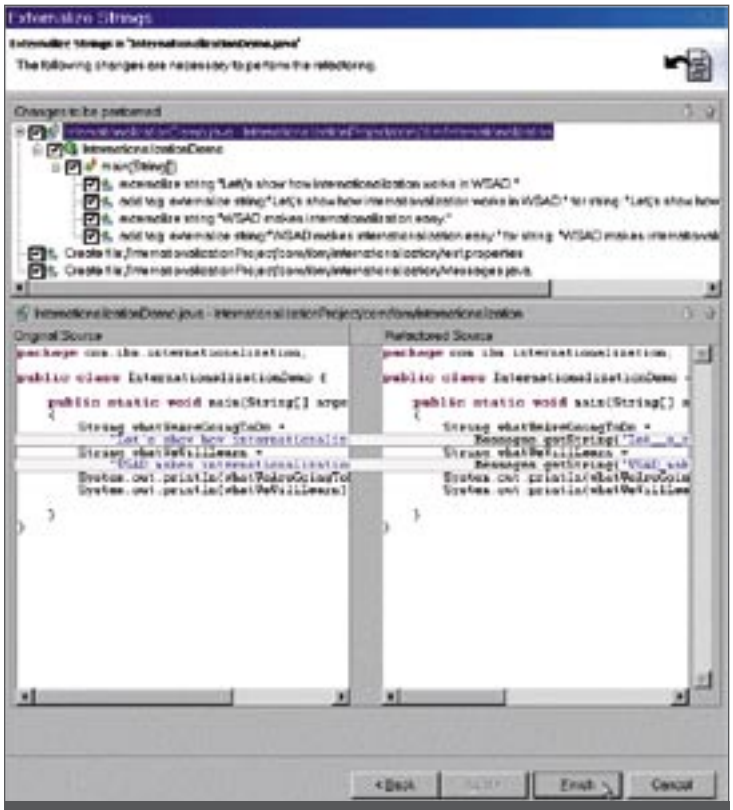

FIG 2: REVIEWING THE PROPOSED CHANGES BEFORE REFACTORING

you are shown precisely how your code will be modified. The original source code is juxtaposed to the proposed refactored source code. When you are done reviewing, click "Finish" to refactor the code. This generates the files Messages.java (see Listing 2) and test.properties (see Listing 3). Messages.java is a resource bundle accessor class that allows you to provide a key and returns the paired String in the test.properties file.

The Old Swap-a-roo

The advantage of the setup we just created is quite powerful. In order to present our application's messages in a different language, all we have to do is simply swap out the test.properties file with one that has the appropriate name/value pairs in our desired language. For example, sticking to our article's code example, we could have an English and a Spanish version of the properties file (see Listing 3).

Conclusion

In our multilingual world, it frequently becomes important to support different languages in our applications. One tedious way to go about doing this might be to create a different version of an application for each language supported. As you can imagine, such an approach can be a nightmare when it comes to code maintenance. As you have seen in this article, WebSphere Studio allows you to support multiple languages in a rather efficient way. Using WebSphere Studio you can extract the existing Strings of your code and have them relocated to a properties file. This properties file can then be swapped out to cater to an audience and their desired language.

Acknowledgment

- The author would like to thank Ricardo Olivieri for his help in making this article possible.

Kulvir Singh Bhogal works as an IBM Software Services for WebSphere consultant, devising and implementing WebSphere-centric solutions at customer sites across the nation. He has over 50 patents pending in a myriad of technology areas.
kbhogal@us.ibm.com

LISTING 1

```
package com.ibm.internationalization;

public class InternationalizationDemo {

  public static void main(String[] args)
  {
        String whatWeAreGoingToDo =
                "Let's show how international-
                ization works in WSAD.";
        String whatWeWillLearn =
                "WSAD makes internationalization
                easy.";
        System.out.println(whatWeAreGoingToDo);
        System.out.println(whatWeWillLearn);

  }
}
```

LISTING 2: MESSAGES.JAVA

```
package com.ibm.internationalization;

import java.util.MissingResourceException;
import java.util.ResourceBundle;

public class Messages {

  private static final String BUNDLE_NAME
        = "com.ibm.internationalization.test";
//$NON-NLS-1$

  private static final ResourceBundle RESOURCE_
  BUNDLE = ResourceBundle.getBundle(BUNDLE_NAME);

  private Messages() {
  }

  public static String getString(String key) {
        try {
                return RESOURCE_
                BUNDLE.getString(key);
        } catch (MissingResourceException e) {
                return '!' + key + '!';
        }
  }
}
```

LISTING 3: TEST.PROPERTIES

```
English version
Let__s_show_how_internationalization_works_in_
WSAD._1=Let\'s show how internationalization
works in WSAD.
WSAD_makes_internationalization_easy._2=WSAD
makes internationalization easy.

Spanish version
Let__s_show_how_internationalization_works_in_
WSAD._1=Veamos como el proceso de internacionali-
zar el software trabaja en WSAD.
 WSAD_makes_internationalization_easy._2=WSAD
hace sencillo el proceso de internacionalizar el
software.
```

WEBSPHERE NEWS

Salesforce.com Announces sforce Toolkit for IBM WebSphere

(New York) – Salesforce.com, a provider of on-demand customer relationship management (CRM) services, has announced the sforce Toolkit for IBM WebSphere. sforce, the world's first on-demand application platform, allows enterprises and solution developers to customize, integrate, and extend salesforce.com to create custom-tailored CRM solutions.

With the new sforce Toolkit for IBM WebSphere, developers and enterprises can leverage WebSphere Application Server and WebSphere Studio Application Developer to build sophisticated CRM solutions that allow key customer and sales information to be integrated across the enterprise – with a fraction of the time and effort required for software-only alternatives. The toolkit is free for both salesforce.com customers and developers using WebSphere.
www.sforce.com

IBM Unveils Software Development Platform Evaluation Packs

(Somers, NY) – IBM has released eight PowerPack evaluation kits on developerWorks to provide software engineers with hands-on experience for evaluating the IBM Software Development Platform. IBM has also launched a new section on its developerWorks site that extends its e-business On Demand Resource Center to help development teams quickly build and deploy software and software services across heterogeneous systems using the IBM Software Development Platform.

Each PowerPack is designed for a specific audience, including the IT analyst, software architect, IT manager, tester, and project manger, and is composed of resources relative to a user's development needs. The PowerPack tracks will enable developers to review and evaluate the IBM Software Development Platform's comprehensive set of products and shared set of proven best practices.

Once a user registers for a PowerPack and selects a track of interest, IBM will send a DVD and a series of e-mails loaded with demos, whitepapers, technical content, and code to help the registrant evaluate software tools from the IBM Software Development Platform. The individual tracks within a PowerPack feature three customized educational programs, each building upon the others to provide a comprehensive technology overview based on a customer's interest.
www-106.ibm.com/developerworks/platform

H&W Releases Performance Scouts Data Collection Agents

(Boise, ID) – H&W has announced the release of DiagnoSys Performance Scouts, the mainframe data collection agents for its DiagnoSys intelligent performance management solution. Together, DiagnoSys and Performance Scouts ensure that crucial Web applications perform optimally, keeping users satisfied and revenue flowing, the company says.

Performance Scouts gather comprehensive metrics from these important mainframe resources: IBM OS/390 or z/OS, TCP, DB2, IBM Transaction Server (CICS), IBM WebSphere for OS/390 and z/OS, WebSphere MQ (MQ Series), and the OS/390 HTTP server. Performance Scouts feed this data into the DiagnoSys intelligent analysis engine, where it is statistically correlated with data from the application resources in other tiers, such as application servers, Web servers, and databases. DiagnoSys analyzes the enterprise-wide components as a single system to give a complete, true view of application performance. It then provides essential information, alerts, and advice to resolve problems proactively.
http://hwcs.com

Agile Combines Product Life Cycle Management with WebSphere Portal

(San Jose, CA) – Agile Software Corporation, a provider of product

IBM OFFERS NEW WEBSPHERE PORTAL V5 CERTIFICATION EXAMS

(Somers, NY) – IBM has announced two new certification exams to help J2EE developers easily update their skills on WebSphere Portal v5. The new exams are designed to help developers become certified faster and at half the cost of competitive certification offerings. In an effort to simplify the certification process and meet marketplace demand for current skills, IBM is offering new certification exams to help developers learn how to maximize the latest features of WebSphere Portal v5 to advance their careers and gain a faster return on their skills investments.

The new WebSphere Portal v5 certification exams are:

- IBM Certified System Administrator - WebSphere Portal v5.0: Test 347, IBM WebSphere Portal v5.0, Deployment and Administration
- IBM Certified Solution Developer - WebSphere Portal v5.0: Test 348, IBM WebSphere Portal v5.0, Application Development

These certification exams are now generally available worldwide and are accessible at www-1.ibm.com/certify/news/20040209.shtml. Information for developers is also available at that site.

life cycle management solutions, has announced Agile 9 platform support for IBM's WebSphere Portal. Agile will deliver WebSphere Portal as an option for the Agile 9 PLM platform to help customers improve employee productivity across the product life cycle. Additionally, customers already using WebSphere Portal will have the ability to incorporate Agile 9 portlets into their existing portal environments.

Agile 9 portlets deployed on WebSphere Portal provide a broad reach of product record information to users within an enterprise and to partners. Deploying portlets increases the usability for end users by empowering them to personalize the look and feel, and customize the way they view information in Agile. Agile portlets can now be deployed with information from other enterprise systems such as ERP, enabling a single dashboard view into related product record business processes.
www.agile.com

200 PARTNERS SIGN ON TO IBM'S ISV ADVANTAGE INITIATIVE FOR SMB IN FIRST YEAR

(Somers, NY) – IBM has announced that 200 independent software vendors (ISVs) have joined its ISV Advantage Initiative in the program's first year. ISVs dedicated to providing open solutions to customers across all vertical industries have committed to go to market with IBM for a majority of their business in the $300 billion small–medium business (SMB) technology market. This includes many companies that have migrated from proprietary offerings to Linux and IBM's Express middleware for the SMB market.

ISV Advantage provides ISVs with comprehensive technical, marketing, and sales support to help meet the specific needs of medium-size businesses. ISVs that enable their applications on IBM's industry-optimized infrastructure software can provide customers with flexible and secure solutions that span multiple computing platforms, including Linux. Currently, more than 70 percent of ISV Advantage participants are actively supporting IBM's software running on Linux.

Unlike other companies that develop applications that compete directly with their partners' offerings, IBM collaborates with ISVs to deliver solutions designed to meet specific customer requirements. AlphaNova, a London-based CRM vendor and new ISV Advantage partner, recently migrated from Microsoft technology to IBM.

"We moved away from the Microsoft environment because of its ever-increasing cost to deploy, which makes it prohibitive for SMBs. CRM requires universal access by almost all employees and partners of an organization, and Microsoft's closed and holistic attitude to use all their products makes cost of ownership very high," said Andros Papageorgiou, CEO, AlphaNova.

"We chose J2EE and WebSphere to provide our SMB customers openness, choice, and lower cost of ownership. IBM's commitment to Linux, which is strategic for us, and the completeness and robustness of the WebSphere platform were other important factors we took into account."
www.ibm.com/partnerworld

Wily Launches Portal Manager 4.0

(Brisbane, CA) – Wily Technology, a provider of Enterprise Application Management, has announced Wily Portal Manager 4.0 for IBM WebSphere Portal 5.0, a first-to-market portal management solution and the only enterprise portal manager to exclusively support IBM WebSphere Portal customers.

The launch of Wily Portal Manager 4.0 builds on the company's strategic direction to manage the performance and availability of next-generation application platform suites that are being implemented by leading enterprises today.

With deep, internal visibility into the portal environment and external visibility into connections to back-end systems, Wily Portal Manager makes it possible for administrators to immediately pinpoint the root cause of performance problems, according to Wily. By providing a unique comprehensive view into each internal portal workflow process and all of its connected systems, administrators can address issues before they impact users, and stay on top of service-level agreements and quality-of-service goals.

"Wily is closely aligned with IBM's product strategy and is committed to its strategic partnership with IBM to provide the highest levels of performance and management solutions for joint customers," said Dick Williams, president and CEO of Wily Technology, Inc.
www.wilytech.com

PalmSource Expands Java Support With IBM WebSphere

(San Jose, CA) – PalmSource, Inc., a provider of Palm OS, an operating system powering next-generation mobile devices and smartphones, has announced it has licensed IBM's WebSphere MicroEnvironment (WME) Java 2 Micro Edition certified runtime environment and WebSphere Studio Device Developer (WSDD) toolset for integration into the Palm OS platform.

Open standards–based software development environments such as WME are expected to create new market opportunities for the more than 275,000 Palm OS developers and to broaden the reach of the Palm OS platform into the growing Java developer community. In addition to native Palm OS applications, WME provides Palm OS developers the tools needed to execute Java MIDlet applications designed to run on Palm Powered smart mobile devices.
www.palmsource.com

RSS FEED NOW AVAILABLE ON WJ HOME PAGE

The WJ home page (www.sys-con.com/websphere) now offers an RSS feed (www.sys-con.com/websphere/feed.rss) to help you keep current with the latest goings-on in the WebSphere community. Sign up today to start receiving timely updates of product announcements, news items, and other important happenings in the world of WebSphere.

Managing J2EE Applications and Their Underlying Infrastructure

BY JIM MCQUAID

The development of production Java applications for deployment on WebSphere or any application server clearly demands a high level of ability. But management of such an application, once deployed, should not require the same level of skill. Operations cannot scale to the demands they face if Java programming skills are required for basic monitoring and management. Consequently, such an environment would present inherently challenging management problems. With today's cost pressures on IT, there is a huge squeeze on Operations to ensure they can monitor and manage the performance and availability of proliferating J2EE-based applications.

Historically, computing systems and applications were predictable, but beginning with client/server and accelerating with the Web-based application server model, applications now have "moving parts" distributed among multiple CPUs, computers, application server containers, and databases. Data can be stored, accessed, and processed by many applications across the network. Browsers, switch fabrics, security policies, and application servers themselves contribute to this complexity. Regardless of this, Operations remains ultimately responsible for performance and availability of the total service.

Needless to say, Operations teams must be equipped to identify and isolate both application and systems problems. Systems experts need to evaluate and resolve the behavior of distributed transactions – down to pinpointing which EJB might be the bottleneck. Traditional systems management frameworks fall short because of their excessive cost, time to deploy, and the consulting services required to become effective. These solutions are monolithic, inflexible, and not well suited to addressing the unique nature of WebSphere and J2EE. Point solutions, while providing greater WebSphere and J2EE management depth, lack the breadth needed to monitor all the key moving parts.

At the end of the day, application servers such as IBM WebSphere provide critical functionality and access to back-end enterprise systems, making their management business critical. It's vital that operational efficiency be achieved with low overhead while preserving the flexibility to monitor virtually any component of the running application.

Seeking to balance the real constraints of Operations groups with the need to manage application servers and J2EE applications themselves, we have developed the concept of five layers of management.

- ***First, is it alive?*** The primary concern of any systems management job is uptime. Aspects of monitoring that address this include: gross load and resource consumption on the computer, simple J2EE server health indicators, and J2EE subsystem health indicators.
- ***Second, core performance metrics.*** In any Java environment, the performance of the JVM (Java Virtual Machine) is critical. The overall load must be within reasonable bounds, and when it's not, you have to know immediately. CPU load, thread count, and memory management are the key performance indicators for this layer.
- ***Third, detailed operations monitoring.*** Here Operations looks for finer-grained information about server performance, communication between clustered servers, unsuccessful login attempts, the JDBC connection pool, EJB pool utilization, and current and leaked connections for the JCA connection pool. A handful of monitoring tasks, returning alerts and/or data every 5 to 10 minutes, can monitor performance down through all the key subsystems.
- ***Fourth, reactive drilldown on performance details.*** When specific performance questions are raised – whether narrow (e.g., one EJB implementing a specific piece of business logic) or broad – additional monitoring is needed on a short-term basis. You must have the ability to monitor an individual EJB or servlet, for example, when circumstances and suspicions warrant. This layer of management will add measurable overhead and is therefore best used only on a problem-solving, reactive deployment basis.
- ***Fifth and last, code analysis and rewrite.*** This is the point at which Operations escalates the problem to Development. Having isolated the basic subsystem or even component involved in performance or availability problems, it's time for the Java expert to delve deeply into the code itself.

Application servers – critical though they may be – are only one component of a larger infrastructure. Operations needs a single window through which the entire infrastructure may be monitored, managed, and secured. But equally important, Operations groups cannot replace developers, nor can they cost-effectively scale to handle their critical mission if they must have the skill sets of a developer.

In Operations, time to value and time to manage is everything. Without a clear understanding of the difference between the missions of the Java developer and the IT operations organizations that must make it all run 24/7, it is easy to saddle one group or the other with a toolset that will not support the real-world demands of each.

Jim McQuaid, director of Monitoring Solutions for NetIQ Corporation, leads the definition and delivery of networked application performance management products for the company. Coauthor of RFC 1944/2544, Benchmarking Methodologies for Network Interconnect Devices, he is a University of Michigan graduate and has held positions at NetScout Systems, Bay Networks, and Analog Devices. Jim has spoken at NetWorld+Interop, COMNET, HP OpenView Forum, TeleCon East, AFCOM, DVC, and Network World Town Meeting. jim.mcquaid@netiq.com

Forget something?

Post-launch is NOT the time to be verifying web applications.

The wild blue yonder of operational monitoring and management is extremely unforgiving. Which means that going live with the monitoring software you used in development is a great way to go dead—quickly! You simply can't support operations if your staff is drowning in details provided by development profiling tools and debuggers. **Let NetIQ cover your apps...with AppManager.**

AppManager—the industry's easiest-to-use Systems Management suite—is a proven management system for monitoring J2EE application servers, databases, operating systems and even end-user response time. NetIQ's AppManager monitors ALL application components—not just your server. **NetIQ. Nobody does UNIX better. Nobody.**

Visit us at www.netiq.com/solutions/web to learn how we can help you address the challenges of your operational monitoring and management.

netiq
Work Smarter.