

No. 1 *i*-Technology Magazine in the World

JDJ

FEBRUARY 2005 VOLUME:10 ISSUE:2



.NET

with Your

Java?



PLUS...

- ▶ Migrating Enterprise Apps Between J2EE App Servers
- ▶ Web Services: Why Can't We Just Talk?

RETAILERS PLEASE DISPLAY UNTIL APRIL 30, 2005

\$5.99US \$6.99CAN



0 09281 01751 6

SEE SHOW BROCHURE INSIDE!



- Join Over **3,000** Developers
- App Server Shoot-Out!
- **FREE** J2EE Tutorial
- JDJ Partner Pavilion
- Seminars and Case Studies

\$1,000,000 Software Giveaway

PAGES 51-60
February 15-17, 2005
Hynes Convention Center
Boston, MA



You deploy to multiple platforms.
You need one solution.



InstallAnywhere

Powerful Multiplatform Deployment



- One Installer Project for Multiple Platforms and Multiple Languages
- Support for Windows, Linux, Solaris and more than 20 other platforms
- Optimized for complex enterprise deployment
- Download today at ZeroG.com

ZERO G
SOFTWARE

Open Source, Open Sesame?



Jeremy Geelan



Editorial Board
 Desktop Java Editor: **Joe Winchester**
 Core and Internals Editor: **Calvin Austin**
 Contributing Editor: **Ajit Sagar**
 Contributing Editor: **Yakov Fain**
 Contributing Editor: **Bill Roth**
 Contributing Editor: **Bill Dudney**
 Contributing Editor: **Michael Yuan**
 Founding Editor: **Sean Rhody**

Production
 Production Consultant: **Jim Morgan**
 Associate Art Director: **Tami Lima**
 Executive Editor: **Nancy Valentine**
 Associate Editors: **Jamie Matusow**
 Gail Schultz
 Assistant Editors: **Natalie Charters**
 Seta Papazian
 Online Editor: **Martin Wezdecki**
 Research Editor: **Bahadir Karuv, PhD**

Writers in This Issue
 Calvin Austin, Paulo Caroli, Bruno Collet,
 Jeremy Geelan, Eric Giguère, Michael Havey,
 Alexander Krappf, Bahar Limaye, Warren MacEvoy,
 Peter V. Mikhailenko, Dave Parsons, Ken Sipe,
 Susan Vickers, Joe Winchester, Michael Juntao Yuan

To submit a proposal for an article, go to
<http://grids.sys-con.com/proposal>

Subscriptions
 For subscriptions and requests for back orders, please send your letters to Subscription Department:

888 303-5282
 201 802-3012
subscribe@sys-con.com

Cover Price: \$5.99/issue. Domestic: \$69.99/yr. (12 Issues)
 Canada/Mexico: \$99.99/yr. Overseas: \$99.99/yr. (U.S. Banks or
 Money Orders) Back Issues: \$10/ea. International \$15/ea.

Editorial Offices
 SYS-CON Media, 135 Chestnut Ridge Rd., Montvale, NJ 07645
 Telephone: 201 802-3000 Fax: 201 782-9638

Java Developer's Journal (ISSN#1087-6944) is published monthly (12 times a year) for \$69.99 by SYS-CON Publications, Inc., 135 Chestnut Ridge Road, Montvale, NJ 07645. Periodicals postage rates are paid at Montvale, NJ 07645 and additional mailing offices. Postmaster: Send address changes to: Java Developer's Journal, SYS-CON Publications, Inc., 135 Chestnut Ridge Road, Montvale, NJ 07645.

©Copyright
 Copyright © 2005 by SYS-CON Publications, Inc. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy or any information storage and retrieval system, without written permission. For promotional reprints, contact reprint coordinator Kristin Kuhnle, kristin@sys-con.com. SYS-CON Media and SYS-CON Publications, Inc., reserve the right to revise, republish and authorize its readers to use the articles submitted for publication.

Worldwide Newsstand Distribution
 Curtis Circulation Company, New Milford, NJ
 For List Rental Information:
 Kevin Collopy: 845 731-2684, kevin.collopy@edithroman.com
 Frank Cipolla: 845 731-3832, frank.cipolla@epostdirect.com

Newsstand Distribution Consultant
 Brian J. Gregory/Gregory Associates/W.R.D.S.
 732 607-9941, BJGAssociates@cs.com

Java and Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries. SYS-CON Publications, Inc., is independent of Sun Microsystems, Inc. All brand and product names used on these pages are trade names, service marks or trademarks of their respective companies.



As Sun open-sources Solaris, and another software development “community” is tugged into being around it, critics are saying – Red Hat’s general counsel Mark Webbink in particular – that the strategy will fail.

“It is much harder to build a community around pre-existing software than one might believe,” Webbink declared, in a statement to the press. (Red Hat and Sun are particularly fierce rivals just now, as anyone knows who has not been asleep in a cave for the past 12 months.)


Sun on the other hand is adamant about the importance of such a community. “The opensolaris.org Web site will be the center for OpenSolaris activity,” it writes on the site’s main page, adding: “Community involvement and input is critical to the success of the OpenSolaris program. If you want us to update you on major events such as the availability of buildable Solaris source code, please provide us with your e-mail address in the text box on the right.”

It set me thinking whether “community involvement” is actually the philosopher’s stone that everyone seems to think it is. Sometimes I find myself wondering, for example, what those who had the idea of naming the “Java Community Process” (JCP) were smoking. How much better it might have been, in hindsight, to have called the JCP just the JC, the Java Council – or some such name that allowed it to achieve its aim of cultivating a wider participation in the specification of new and existing Java APIs, the “process” part, without burdening it with the additional notion of “community” which to this day seems to be as much a bane as a boon.

Unlike that other JCP, the Japanese Communist Party, which was founded in 1922, the Java Community Process was founded only in 1998 and has guided the development and approval of Java technical specifications since then. James Gosling was quoted in 1999 as saying the JCP was all about “getting the managers

out of the room and letting the technical people do their thing.” A noble aim, as any developer would agree; though it’s by no means clear that managers are without a role to play, and many are indeed active in the JCP.

In last month’s JDJ David Parsons, Ilan Kirsh, and Mark Cranshaw looked at the Java Technology for the Wireless Industry specification, a recent example of the JCP seeking to help establish a set of industry-wide goals and expectations for Java-based applications and services through the JSR-185 Java Technology for the Wireless Industry (JTWI). And in this issue we have a full-fledged feature looking at the JCP from various perspectives, not



least that of its current chairman and longtime JDJ contributor and advisor, Onno Kluyt.

Java architect and instructor Allen Holub recently wrote “bad standards are destructive. Just look at created-by-committee junk like EJB and JSF” – and

in the same opinion piece declared: “My outsider view of the Java Community Process is that it doesn’t work.”

But the JCP community is an important place for developers to be involved in. With the JCP version 2.6 in place now, the community gets even more opportunity to influence and contribute to standards while they are being developed.

Probably Churchill’s observations about democracy apply. “Democracy is the worst form of government,” he famously said, adding “except for all those others that have been tried.” As with democracy, so with the JCP. And so, too, with OpenSolaris.org. While everything could always be better, things could be an awful lot worse.

So if a community – whether it is the JCP, OpenSolaris.org, or anything else – isn’t the way you’d like it to be, first join it; then change it. As the playwright Henrik Ibsen once wrote: “A community is like a ship, everyone ought to be prepared to take the helm.”

Jeremy Geelan is group publisher of SYS-CON Media and is responsible for the development of new titles and technology portals for the firm. He regularly represents SYS-CON at conferences and trade shows, speaking to technology audiences both in North America and overseas.

jeremy@sys-con.com

Recognize real value

Check out Authentic® 2005,
the genuinely FREE XML and
database content editor from Altova®.

Detected in Version 2005:

- Relational database content editing
- XPath 2.0 support
- XSLT 2.0 support

Put Authentic 2005 into action as the non-technical data entry element of advanced XML-based document frameworks and management oriented databases. Deploy Authentic 2005 to enable business personnel to edit content without being exposed to the underlying technology. They simply key in electronic forms using the word-processor style interface and their data is instantly adapted behind the scenes. Get your technology undercover!

Download Authentic® 2005 today:

www.altova.com

JDJ contents

JDJ Cover Story

42

...the direct way

Hosting ActiveX/.NET Controls in a Java Application...



by Stanley Wang

Features

36



Web Services and XML Special

by Onno Kluyt

FROM THE GROUP PUBLISHER

Open Source, Open Sesame?

by Jeremy Geelan

.....3

METHODOLOGIES

Why Use Extreme Programming?

Not just another development methodology

by Troy Holmes

.....16

DESKTOP JAVA VIEWPOINT

Software Testing Shouldn't Be Rocket Science

by Joe Winchester

.....40

JAVA ENTERPRISE VIEWPOINT

Outsourcing to Students

by Yakov Fain

.....6

CORE AND INTERNALS VIEWPOINT

**Web Services:
Why Can't We Just Talk?**

by Calvin Austin

.....28

LABS

Oracle Application Server 10g

Reviewed by Mark Euler

.....48

ENTERPRISE

**Design a Reference Data Service
in J2EE Application Server Clusters**

Improve the overall system performance

by Di Li

.....8

TECHNOLOGY

**Opening Java Projects up to
Open Source**

*The open source community has
room enough for all*

by Timothy Fisher

.....30

JSR WATCH

**From Within the Java Community
Process Program**

Java technology interfaces for Web services

by Onno Kluyt

.....62



24

**Migrating Enterprise Applications
Between J2EE Application Servers**

by Ajit Sagar



Yakov Fain
Contributing Editor

Outsourcing to Students

One of my resolutions this year is to start teaching part-time Java-related classes in some college. That's why I started browsing the computer science course lists that are being offered this year. While graduate-level programs offer many interesting courses, the situation is different in the undergrad world.

Some schools keep teaching how to multiply matrices in Ada or work with algebraic expressions in Prolog. Half of the courses are preparing professionals who will be operating on another planet. Information systems programs look a little more down to earth. Meanwhile, many college graduates are having a hard time finding their first jobs because many entry-level programmers jobs are being outsourced overseas, and it'll stay this way as long as it makes financial sense for businesses. Unfortunately, student loans have not been outsourced...

A catch-22 situation with experience makes things even worse: we can't hire you because you don't have the industry experience. How can I get this experience if no one hires me? Some people try to find volunteer programming work to get a foot in the IT door. Some graduates join open source projects, and some job applicants just lie on their résumés.

I have a plan: instead of outsourcing projects to developing countries, businesses should offer them to the local colleges. I'm not talking about simple pilot or proof-of-concept projects; I mean the real ones. This plan requires commitment and the cooperation of academia and businesses. These are some thoughts that come to mind:

- Colleges have to include more classes on software engineering and modern technologies in the undergraduate programs. Here are some of the candidates: Application Servers, Service-Oriented Architecture, Design and Development of J2EE Applications, Applying Design Patterns, Data Modeling, Business Intelligence, and UML.

- Colleges form teams of programmers starting from the students' junior year. Faculty members lead these teams. Information about these teams (résumés, previous projects, GPAs) has to be published on the Internet and be publicly available, and businesses need to publish their project descriptions so student teams can bid on these projects.
- Colleges make their labs, networks, and support personnel available for the teams. If needed, businesses can lease additional hardware to the college for the duration of the project.
- Most of the students study Java programming during their freshman and sophomore years. Many Java components are available for free or through open source licenses: IDE, version control systems, project build tools, bug reporting systems, application servers, etc. Businesses will purchase any additional required software for a fraction of the cost using heavily discounted academic prices.
- Business managers pick and interview teams for their projects based on the college reputation, available skill sets, location, and other criteria.
- Business lawyers prepare a contract with a selected team that defines the obligations of each party, deliverables, cost of development, and potential penalties.
- The turnover rate is usually high on the projects that are outsourced to developing countries, which won't be the case with student teams. On the other hand, there is a risk of not having developers during midterms and final exams. However, since the exam schedules are known in advance, the "freeze time" can be planned accordingly.
- Most of the business managers dealing with developers from other countries complain that cultural differences are a huge problem. Guess what? This won't be a problem if you outsource the project to students who live in the same country and speak your language.

—continued on page 28

Yakov Fain works as a Java architect for a major bank in New York City. He wrote the book *The Java Tutorial for the Real World*; an e-book *Java Programming for Kids, Parents and Grandparents*; and several chapters for the book *Java 2 Enterprise Edition 1.4 Bible*. Yakov holds a masters degree in applied mathematics. For more information please visit www.smartdata-processing.com.
yakovfain@sys-con.com

President and CEO:
Fuat Kircaali fuat@sys-con.com
Vice President, Business Development:
Grisha Davida grisha@sys-con.com
Group Publisher:
Jeremy Geelan jeremy@sys-con.com

Advertising

Senior Vice President, Sales and Marketing:
Carmen Gonzalez carmen@sys-con.com
Vice President, Sales and Marketing:
Miles Silverman miles@sys-con.com
Advertising Sales Director:
Robyn Forma robyn@sys-con.com
National Sales and Marketing Manager:
Dennis Leavey dennis@sys-con.com
Advertising Sales Managers:
Megan Mussa megan@sys-con.com
Kristin Kuhnle kristin@sys-con.com
Associate Sales Managers:
Dorothy Gil dorothy@sys-con.com
Kim Hughes kim@sys-con.com

Editorial

Executive Editor:
Nancy Valentine nancy@sys-con.com
Associate Editors:
Jamie Matusow jamie@sys-con.com
Gail Schultz gail@sys-con.com
Assistant Editors:
Natalie Charters natalie@sys-con.com
Seta Papazian seta@sys-con.com
Online Editor:
Martin Wezdecki martin@sys-con.com

Production

Production Consultant:
Jim Morgan jim@sys-con.com
Lead Designer:
Tami Lima tami@sys-con.com
Art Director:
Alex Botero alex@sys-con.com
Associate Art Directors:
Abraham Addo abraham@sys-con.com
Louis F. Cuffari louis@sys-con.com
Richard Silverberg richards@sys-con.com
Assistant Art Director:
Andrea Boden andrea@sys-con.com

Web Services

Information Systems Consultant:
Robert Diamond robert@sys-con.com
Web Designers:
Stephen Kilmurray stephen@sys-con.com
Matthew Pollotta matthew@sys-con.com

Accounting

Financial Analyst:
Joan LaRose joan@sys-con.com
Accounts Payable:
Betty White betty@sys-con.com
Accounts Receivable:
Steve Michelin smichelin@sys-con.com

SYS-CON Events

President, SYS-CON Events:
Grisha Davida grisha@sys-con.com
National Sales Manager:
Jim Hanchrow jimh@sys-con.com

Customer Relations

Circulation Service Coordinators:
Edna Earle Russell edna@sys-con.com
Linda Lipton linda@sys-con.com
Monique Floyd monique@sys-con.com
JDJ Store Manager:
Brunilda Staropoli bruni@sys-con.com

WebRenderer™



Standards compliant embeddable web browser component

WebRenderer is a cutting edge embeddable Java™ web content rendering component that provides Java applications and applets with a fast, standards compliant HTML and multimedia rendering engine. WebRenderer provides a feature-packed API including complete browser control, a full array of events, JavaScript interface, DOM access, document history and more.

Why WebRenderer?

- Standards Support (HTML 4.01, CSS 1 & 2, SSL, JavaScript, XSL, XSLT etc.)
- Exceptionally Fast Rendering
- Predictable Rendering
- Scalability (deploy in Applications or Applets)
- Security (based on industry standard components)
- Stability and Robustness

Embed WebRenderer to provide your Java® application with standards compliant web content rendering support.

To download a 30 day trial of WebRenderer visit
www.webrenderer.com

JadeLiquid™
Software

Copyright JadeLiquid Software 2004. JadeLiquid and WebRenderer are trademarks of JadeLiquid Software in the United States and other countries. Sun, Sun Microsystems, the Sun Logo and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All other trademarks and product names are property of their respective owners.

Design a Reference Data Service in J2EE Application Server Clusters

by Di Li

Improve the overall system performance

A multitiered J2EE-based enterprise application is usually deployed on multiple J2EE application servers running on geographically separated machines. It may also integrate the services distributed in heterogeneous applications within an enterprise (such as an ERP application or a mainframe legacy application). The Java application servers may also be grouped in clusters for scalability and high availability. Improving overall system performance is becoming a challenging task.

To build a high-performance J2EE application, we will create a reference data service by downloading frequently used data from a mainframe legacy system to each J2EE application server, and design an efficient object repository to control object access and data synchronization with component-based analysis and design, the publish/subscribe messaging model, the Java event model, the Immutable Object pattern, and the Strategy pattern.

Before we discuss any performance tuning strategies, let's take a closer look at J2EE-based applications. A typical J2EE application consists of Web modules deployed on the presentation tier, EJB modules deployed on the business logic tier, and connector modules deployed on an integration tier and EIS tier. Figure 1 portrays a sample collaboration of components in a typical J2EE application.

The most time-consuming processes of the collaborations are the communications between the connector components and mainframe legacy application, the JDBC accesses between the data access components and database, and the remote method invocations between the Web components and the EJB components.

We can categorize business entities in a business domain into two groups: the static group that includes the entities whose states can only be changed by end users, and the dynamic group that includes the entities whose states may be updated through the business process. A sample of the static group is a profile of a customer in a Web-based banking system: the address of a customer can be updated by a bank teller. A sample of the dynamic group is an account of a customer: the balance of an account may be changed through the deposit, transfer, payment, or withdrawal process.

A reference data service is a façade that internally calls a mainframe legacy system that handles data for all applications in an enterprise. This article discusses an intelligent solution for a reference data service: create a cluster-wide object repository to store the data changed only by end users (usually saved in a mainframe legacy system) in the memory of each J2EE application server. The following is a summary of the benefits offered by this solution:

1. Automated vertical sync – the states of each object in any object repository are synchronized with the underlying mainframe legacy system.
2. Automated horizontal sync – the states of each object in different object repositories are kept in sync.
3. Data access service – provides coarse-grained interfaces to retrieve objects and their relationships from the object repository. Shields the complexity of location, access, transport, and conversion.
4. When the state of an object in the legacy system is changed by an end user, the change is immediately represented at each object repository.
5. Improving overall system performance by reducing the remote method invocations between the Web components and the EJB components, and the communications between the connector components and mainframe legacy system.

Related Techniques

Publish/Subscribe Messaging Model

The Publish/Subscribe messaging model (see Figure 2) enables an application to send messages to a topic, and multiple subscribers to that topic receive all the messages sent to that topic.

In the reference data service solution, the in-memory object repository distributes to multiple J2EE servers. These servers are grouped into different clusters. Some servers may be out



Di Li, a Sun Certified

Enterprise Architect, is an expert in the area of business process modeling, enterprise architecture design, and software architecting. He has over 10 years of experience in software engineering and technical architecture. He worked in providing the end-to-end solution of financial services, e-business, and enterprise application integration.

victor1998@hotmail.com

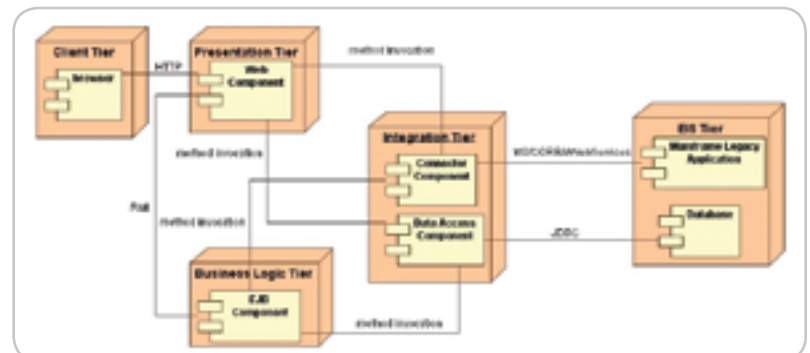


Figure 1 A sample collaboration of components in a typical J2EE application

Remember the good old days?



Ah, yes. Those were simpler times. But you're not nostalgic for old technology. If your Java application relies on SQL Server data, your database driver should give you the performance, reliability, and functionality Microsoft® SQL Server deserves. **DataDirect presents today's JDBC** – the SPECjAppServer/ECperf leader. Featuring Windows Authentication support, J2EE certification, and advanced 3.0 specification compliance for SQL Server 2000 and SQL Server 7.

Don't use yesterday's technology.
Get current with **DataDirect Connect™ for JDBC.**

www.datadirect.com/jdj
(800)-876-3101

DataDirect™
TECHNOLOGIES

of service and leave the clusters. Other servers may join the clusters and provide services. To keep data in sync among the servers, we use MQ JMS Pub/Sub messaging as the delivery infrastructure of the event. It decouples the destination of an event from the sender. The number of message producers and consumers can be dynamically changed. So the repository in one server is not aware of the existence of the repository in another server. The startup or the shutdown of one server will not impact the service of the whole object repository.

Java Event Model

The traditional solution of a refresh mechanism is based on a timer: the data in the memory is reloaded from a legacy system only when the timer expires. Any change that occurs before the timer expires is not represented in the memory. The timer-based solution may also cause system overhead when the timer expires but the data is never actually updated.

My refresh mechanism solution uses the Java event model (see Figure 3), which is based on the concepts of the event, event source, and event listener. The event source object maintains a list of listeners that are interested in receiving events that the source produces. When the Event Source object generates an event, it notifies all listeners that an event has occurred. The underlying

mainframe system works as the event source, and the in-memory repository is the event listener. The event-based refresh solution updates the repository only when the state of an object is changed in the mainframe system.

Strategy Design Pattern

The Strategy Design pattern defines a family of algorithms, encapsulates each one, and makes them interchangeable. Strategy lets the algorithm vary independently among clients that use it.

In my solution for reference data service, the Strategy pattern is used to fulfill the task of initializing the object repository in each J2EE server. As shown in Figure 4, the RepositoryController component in the business-logic tier delegates that task to the MQConnector component in the integration tier. The MQConnector fetches data from a legacy application and returns the result to the RepositoryController, which completes the task of loading data from a mainframe application to its in-memory repository by calling its strategy (MQConnector).

However, the RepositoryController component in the presentation tier sends that task to the ModelProxy component. The ModelProxy calls the remote RepositoryController in the business-logic tier to retrieve the objects from its repository, and then returns the result to the RepositoryController in the presentation tier. So the RepositoryController in the presentation tier uses its strategy (ModelProxy) to initialize its repository by getting the data from the RepositoryController in the business tier rather than from the mainframe application.

Immutable Object Pattern

The Immutable Object pattern increases the robustness of objects that share references to the object and reduces the overhead of concurrence access to the object. It accomplishes this by not allowing an object's state information to change after it's constructed.

We use the Immutable Object pattern to achieve the data integrity of objects stored in repositories. One object can be read by any component, but can only be updated or deleted by the RepositoryController class. The RepositoryController class implements both the feature of the vertical sync, which

ensures that the state of the objects in a repository is synchronized with the underlying mainframe legacy system, and the feature of the horizontal sync, which ensures that the objects are safely kept in sync in each J2EE application server.

The Use Case:

Employee Relationship Management (ERM)

Our example is based on a Web-based ERM System. There are two typical business entities in this domain: Project and Employee. These two entities have many-to-many relationships: a project is developed by many employees, and an employee may join one or many projects. Only the end users can change the states or relationships of these two entities. This article illustrates that by storing the states and relationships of the Project and Employee entities in the in-memory object repository, the reference data service solution not only handles the state change of the business entities, but also their relationship change.

Since the in-memory copy of the Project entity and the Employee entity exist on each J2EE application server (usually saved in the mainframe legacy system), any component deployed in the J2EE application server cluster (such as a servlet component in a Web module) can retrieve the states or relationships of the Project and Employee from the local in-memory repository rather than from the remote mainframe legacy system.

For demonstration purposes, in this article the solution for the reference data service is built on an IBM WebSphere Application Server and WebSphere MQ. But it could be any J2EE application server and message-oriented middleware that support the Java Message Service.

Architectural Model Design

The architecture of the reference data service solution consists of three components: the repository controller, the event handler, and the MQ Communication Controller.

MQ Communication Controller

The MQ Communication Controller is designed for delivering the event of updating an object state to or from the Repository Refresh Topic. It's built on the JMS Publish/Subscribe messaging model.

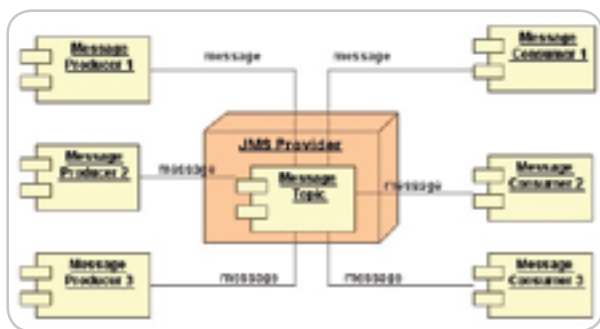


Figure 2 Publish/Subscribe Messaging model




Figure 3 Java event model

Your potential. Our passion.[™]
Microsoft

Feel free to think big.

With Visual Studio[®].NET 2003, you write less code than with Microsoft[®] Visual Basic[®] 6.0, so you can turn that big idea into reality faster than you ever thought possible. You get a new IDE that leverages your existing skills in a tool that gives you more robust code, improved IntelliSense[®] and better deployment. All of which means you're more productive, and more ready to take on your biggest ideas. Find out more at msdn.microsoft.com/visual


Microsoft
Visual Studio

© 2004 Microsoft Corporation. All rights reserved. Microsoft, IntelliSense, Visual Basic, Visual Studio, and the Visual Studio logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Event Handler

The event handler is designed for implementing the horizontal sync feature among J2EE application servers.

As shown in Figure 5, when a repository controller updates, adds, or deletes an Employee object (it could be a Project object) from the repository it contains, it generates an event of updating an object state that wraps the Employee object. The event handler is notified of this event, and then it calls the MQ Communication Controller to publish this event to the Repository Refresh Topic.

When the MQ Communication Controller receives an event from the Repository Refresh Topic, it notifies

the event handler of the event. The event handler first identifies whether the event is published by itself. If not, the event handler retrieves the object wrapped in the event (for example, an Employee object), then sends the object to the repository controller. The repository controller replaces the object in the repository with the received object.

Repository Controller

The repository controller is designed to provide data access service and implement vertical synchronization between the in-memory object repository and mainframe legacy system.

When an end user sends a request to update an object state (such as a Project object) to a J2EE Web server, the servlet component calls the local repository controller in the same Web server. The local repository controller invokes a remote repository controller in one of the J2EE EJB servers by sending a project object. The repository controller in the J2EE EJB server updates the related record in the legacy application, updates its repository with that project object, and then creates an event of updating an object state that wraps the project object. The event handler in the same J2EE EJB server is notified of this event. Finally the repository controller in the Web server updates its object repository with the project object.

After an event handler receives an event of updating an object state, it retrieves the object (for example, a Project object) wrapped in the event, and then sends that object to the repository controller. The repository controller compares the timestamp of the object it receives with the timestamp of the object it contains to identify which is the latest version of the object. (The timestamp of an object is set when that object is created and will never be changed, so each object owns a unique timestamp.) If it's the latest version, the repository controller updates the object stored in the repository.

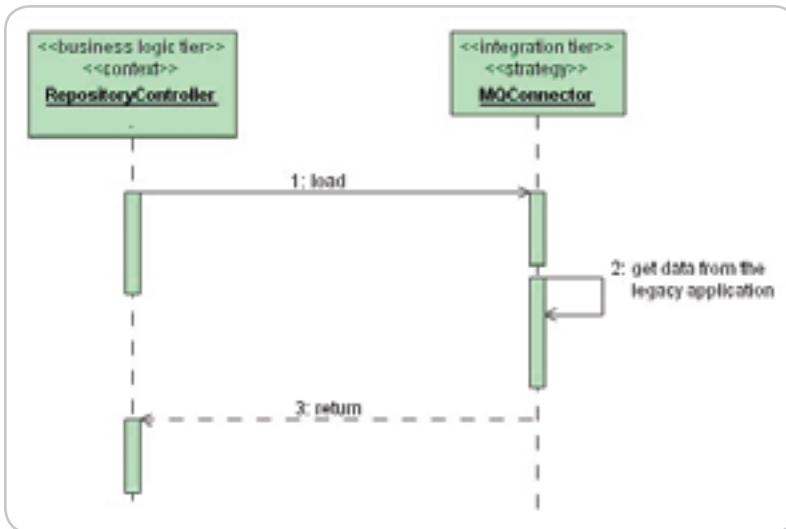


Figure 4 Loading data from a mainframe application to an in-memory repository

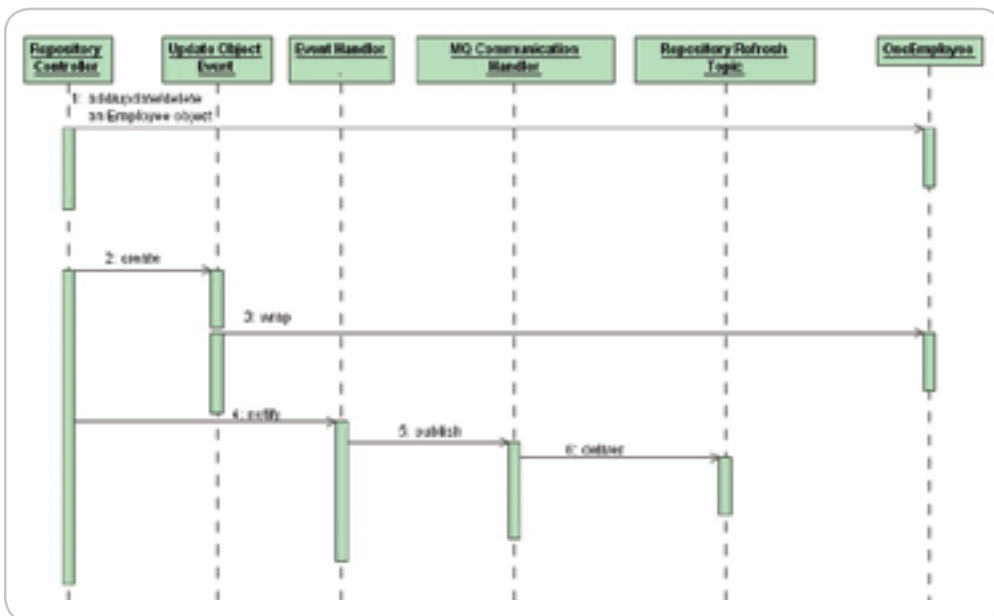


Figure 5 Add/Update/Delete an Employee object from in-memory repository

Solution Model Design

Figure 6 shows the solution model of a reference data service.

Project Class and Employee Class

The Project class and Employee class represent the business entities of project and employee. They implement the Immutable Object pattern. The “isReadyOnly” attribute and “isExpiry” attribute restrict the setter/getter methods. The “timestamp” attribute is set when an object is created. It’s used to identify the version of an object in the system and make certain that only the latest version of an object is stored in each object repository. Listing 1 illustrates how to implement the Immutable Object pattern. Listings 1–5 can be downloaded from www.sys-con.com/java/sourcec.cfm.

The “employeeId” attribute of the Project class and the “projectId” attribute of the Employee class are used to re-

FEBRUARY

BUZZ

Jtest®

An Ounce of Prevention

To maximize the benefits of coding standard analysis, it's critical that you recognize coding standards as a means of preventing errors—not detecting them. Many developers are disappointed if a coding standard violation doesn't point them to an obvious bug. When they explore a violation and find an error-prone construct rather than an error, they think that the coding standards aren't useful, eventually stop investigating violations, and later stop performing coding standard analysis altogether. This speaks to a fundamental problem with the software industry: Most development and testing teams are concerned with *removing* errors, but not with *preventing* them.

Error prevention involves correlating each error to the exact point in the development process that allowed it, then fixing that part of the process. This prevents the need to debug applications after the fact, and produces an exponential increase in product quality. Error prevention is very different than error detection, which is the process of finding and fixing errors *after* an application is built. When development focuses only on error detection, the flawed process that generated those errors is left uncorrected and the errors continue to occur. Furthermore, this misplaced focus has a negative impact on software quality and team productivity.

The coding standards promoted by industry experts were produced in response to other developers' mistakes and development process flaws. The experts performed the most difficult task: figuring out why errors were occurring and how to prevent them. To prevent errors, you just follow the coding standards that the experts designed and tested. Essentially, you get to benefit from the lessons learned from other developers' mistakes without having to suffer the consequences of personally committing those mistakes.

— Adam Kolawa, Ph.D.
Chairman/CEO of Parasoft

Deliver better Java code in less time, with fewer resources.



Parasoft® Jtest® makes Java unit testing and coding standards analysis fast and painless.

With just a click, you can verify that your Java code is robust, high quality and secure. Jtest automatically verifies compliance to hundreds of coding rules while automatically generating and executing JUnit test cases — creating harnesses, stubs and inputs.

See how Jtest can enhance your Java development efforts...

Visit our Parasoft Jtest Resource Center at www.parasoft.com/jtest_JDJ for the latest information — including white papers, webinars, presentations, and more.

For Downloads go to www.parasoft.com/Jtest_JDJ

Email: Jtest_JDJ@parasoft.com Call: 888-305-0041 x2174



Features

- Full Eclipse integration
- Quick Fix automatically corrects errors
- Automatically generates JUnit test cases
- Customizable testing and reporting

Benefits

- Makes error prevention feasible and painless, which brings tremendous quality improvements, cost savings, and productivity increases
- Makes unit testing and coding standard compliance feasible and painless
- Encourages the team to collaborate on error prevention

Platforms

- Windows 2000/XP
- Linux
- Solaris

Contact Info

Parasoft Corporation
101 E. Huntington Dr., 2nd Flr.
Monrovia, CA 91016
www.parasoft.com

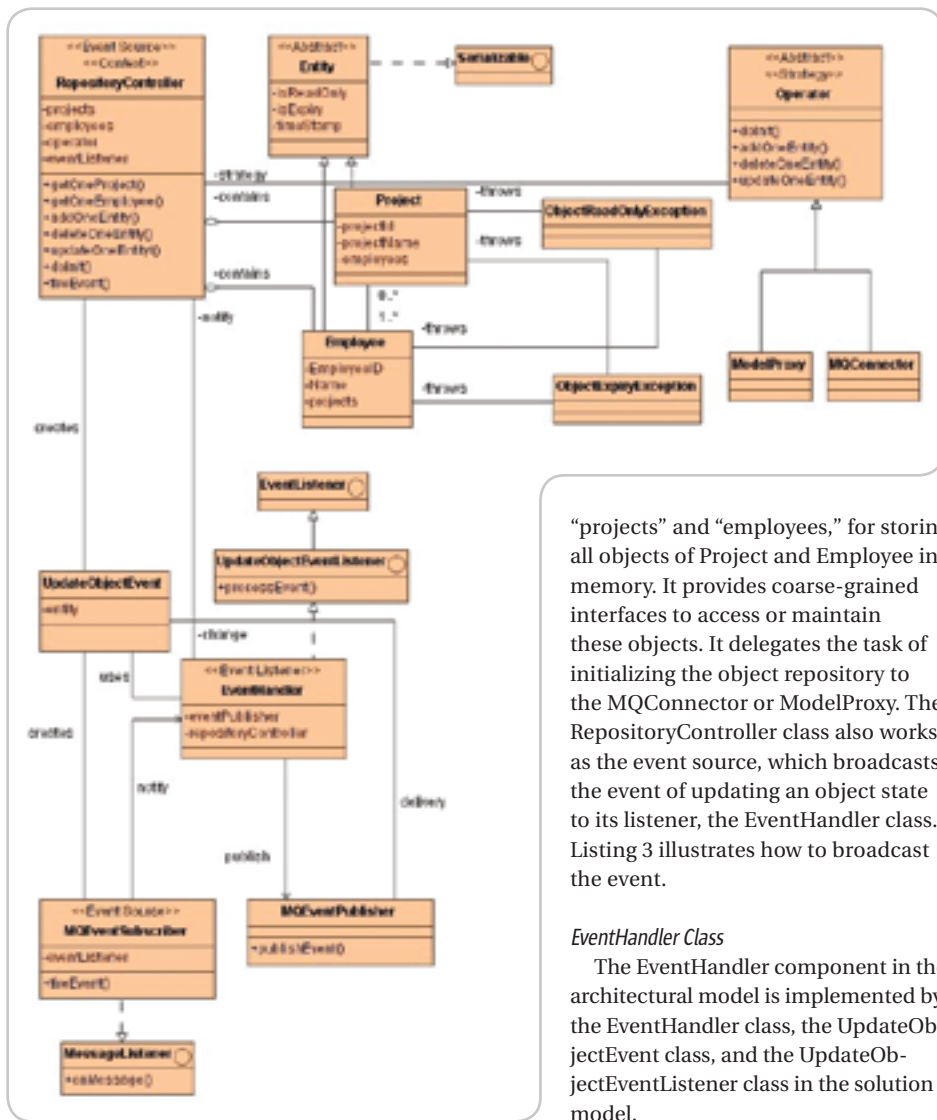


Figure 6 The solution model of the reference data service

cord the relationship between the Project business entity and the Employee business entity. The “employeeIds” attribute saves the IDs of all employees involved in a specific project. On the other side, the “projectIds” attribute records the IDs of all projects that one employee has joined. Listing 2 shows how to retrieve all projects for a specified employee.

Repository Controller Class

The Repository Controller component in the architectural model is implemented by a group of classes in the solution model: the RepositoryController class, the Operator class, the MQConnector class, and the ModelProxy class.

The RepositoryController contains two collection types of attributes,

“projects” and “employees,” for storing all objects of Project and Employee in memory. It provides coarse-grained interfaces to access or maintain these objects. It delegates the task of initializing the object repository to the MQConnector or ModelProxy. The RepositoryController class also works as the event source, which broadcasts the event of updating an object state to its listener, the EventHandler class. Listing 3 illustrates how to broadcast the event.

EventHandler Class

The EventHandler component in the architectural model is implemented by the EventHandler class, the UpdateObjectEvent class, and the UpdateObjectEventListener class in the solution model.

The EventHandler class works as the event listener of the RepositoryController class and the MQEventSubscriber class. If the event source is the RepositoryController, the EventHandler calls the MQEventPublisher to publish an event to the Repository Refresh Topic. If the event source is the MQEventSubscriber, the EventHandler invokes the RepositoryController to update the states of an object in memory. Listing 4 demonstrates how the EventHandler class processes a received event.

MQEventPublisher Class and MQEventSubscriber Class

The MQEventPublisher class and MQEventSubscriber class in the solution model implement the MQ Communication component in the architectural model.

The MQEventPublisher class is used to publish an event to the Repository Refresh Topic. The MQEventSubscriber class is designed for subscribing to an event from the Repository Refresh Topic. When receiving an event, the MQEventSubscriber class first determines whether the event is published by the MQEventPublisher class in the same J2EE application server. If not, it notifies its listener with the event received. Listing 5 demonstrates how the MQEventSubscriber class processes the event received.

Summary

This article described a reference data service solution that can improve the overall system performance of a multitiered J2EE-based enterprise system. The core of the solution is that any components running in the J2EE application server clusters can retrieve the state of the entities from the local object repository rather than from a remote mainframe legacy system. As a result, the overall system performance is improved by reducing the remote method invocations between Web components and EJB components, or by reducing communication between the connector components and the mainframe legacy system. The solution is designed and implemented by combining the following techniques and technologies:

- Design a refresh mechanism with the Java event model.
- Create a mechanism for maintaining data integrity with the Immutable Object pattern.
- Implement a mechanism for initializing the object repository with the Strategy pattern.
- Build an event delivery mechanism with the Publish/Subscribe messaging model.

Resources

- *Java Messaging Service Specification 1.1*: www.java.sun.com/products/jms
- Yusuf, K. (2003). *Enterprise Messaging Using JMS and IBM WebSphere*. Prentice Hall PTR.
- Gamma, E.; Helm, R.; Johnson, R.; and Vlissides, J. (1995). *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley.

Get the complete picture



Features, Performance and Control

Discover the ILOG JViews Graphics Components

You're developing a sophisticated user interface for a desktop, applet or servlet application – it needs to provide displays that go far beyond what Swing and HTML offer. How can you be sure it will have the features, performance, customization and scalability to enable your end-users to make better more informed decisions, faster?

With ILOG JViews, you get comprehensive graphical libraries & tools, resources, and maintenance services so you can focus on the implementation, confidently completing your application in less time and at less cost.

Quickly and easily build:

- Gantt and resource displays
- Graph layouts, diagrams, workflows
- Geographic map displays
- Realtime data charts
- Custom monitoring and control screens
- Network and equipment management screens

**Get a JViews Info Kit – Learn more, test drive an Eval.
Go to: jviews-info-kit.ilog.com or Call: 1-800-for-ILOG**



Changing the rules of business™

Why Use Extreme Programming?

by Troy Holmes

Not just another development methodology

Extrême Programming (XP) has been an accepted form of software development for about eight years now. Many of the concepts found in this lightweight method of development have been implemented into the software shops without even the awareness that they were XP techniques. XP takes many of its fundamentals from other iterative development methodologies, including RAD and JAD. However, unlike those methodologies, which are more of a rapid prototyping approach, XP creates individual components that can be quickly developed and integrated into a larger software system. This article is written with the intent to help software developers understand the key concepts of this methodology and to provide a framework for implementing these processes into a business.

History

XP was created by Kent Beck and Ward Cunningham back in 1996, during a development project at Daimler Chrysler. Beck and Cunningham felt it was time for a better development process because the standard waterfall method left too many gaps and resulted in far too few successful implementations. The fundamental goals of XP were to increase communication, simplify the development process, and obtain feedback from the customer to ensure that requirements were met. XP was one of the first development processes that fell into the realm of iterative programming. This form of development changed the way developers think about the software life cycle because it provided a means to quickly demonstrate requirements with a functional application. Iterative programming both

enables and requires modular design. This process can transform a large development effort into a smaller, task-oriented endeavor. This can be crucial when dealing with a project that must support changing customer requirements. The waterfall method typically requires too much time and results in a completed development project that is outdated by the time it is released. XP was created to give the customer a method by which to prioritize the requirements for software, thereby enabling the best bang-for-the-buck approach.

Key Concepts of XP

This methodology is designed around the typical common sense approach to development; however, these fundamentals are outlined in detail to ensure your understanding and to highlight their importance.


Planning

Planning is the cornerstone of any project's success. During the planning cycle, it is important to communicate with the customer to determine the scope and the schedule of the project. These two factors are used in your typical triangle approach (scope, schedule, resources) and will help to determine which elements have a higher priority. This information will then be used in laying out the release planning. One of the fundamental differences in the planning phase used in XP is the implementation of user stories as a way of capturing use cases. User stories are written by the customer and are

simple requirements that must be met by the software. It is important to note that user stories captured during the planning phase should be high level and should not include any technical jargon. In addition, stories should be small enough to implement in no more than three weeks. These stories should not be longer than a few sentences and should provide enough detail for the developer to estimate the time required to implement. During development, the programmer should ensure the understanding of the user story by communicating his or her version of the implementation with the customer. This information will be used to generate acceptance test code, thus making sure that customer requirements are met.

Designing

The design phase should be a time to validate overall functionality while also focusing on the details of the individual user stories. One of the key concepts to keep in mind during the design phase is to keep it simple. It is extremely important that developers do not add unnecessary complexity during the initial design. XP promotes the idea of refactoring code; refactoring should be an integral process, shared by all developers during development of the software. Therefore, a simple approach is the best approach for initial design. It is also important to outline naming standards during the design phase. At this stage, it will be easier for developers to agree upon and understand the naming conventions. It is important to select names that all team members can relate to without



Troy Holmes works as a senior architect for Prizum Inc. in Reston, VA. He has been working in the IT industry for more than 15 years, with six years primarily focused on the J2EE area. During that time, he has implemented several large-scale enterprise applications. His expertise encompasses all realms of the system development life cycle.

troy.holmes@prizum.com



21 WAYS TO USE SPREADSHEETS IN YOUR JAVA APPLICATIONS

A free offer for readers of *Java Developer's Journal*!

Formula One e.Spreadsheet Engine:

Finally, there's a *supported*, Pure Java tool that merges the power of Excel spreadsheets and Java applications.

- 1 Automatically generate dynamic Excel reports. No more manual querying and cutting-and-pasting to create Excel reports!
- 2 Manage calculations and business rules on J2EE servers with Excel files. No more translating Excel formulas to Java code!
- 3 Embed live, Excel-compatible data grids in applets and Java desktop applications. No more static HTML or presentation-only data grids!

Download this quick-read white paper and trial today!




Download your trial and test our demos and sample code. See for yourself how the Formula One e.Spreadsheet Engine can help your Java application leverage the skills of Excel users in your business.

<http://www.reportingengines.com/download/21ways.jsp>



888-884-8665 • www.reportingengines.com
sales@reportingengines.com

Need to deliver reports from your J2EE application or portal server? Try the Formula One e.Report Engine!



Build reports against JDBC, XML, Java objects, BEA Portal Server logs, BEA Liquid Data, and other sources visually or with Java code. It's embedded! No external report server to set up. Unlimited users and CPUs per license.

<http://www.reportingengines.com/download/f1ere.jsp>

learning detailed business rules. This will ensure a quicker understanding of the overall system.

One of the pillars of the XP design process is collaboration. It is critical that collaboration meetings are held. This is referred to as Use Class, Responsibilities, and Collaboration (CRC) sessions in XP speak. During this time, the developer of an individual component will explain how his or her object will communicate with other objects in the system. This will provide a means to quickly uncover problematic issues and shortcomings in the design. The CRC sessions are written on cards that can quickly be converted into collaboration diagrams if a more permanent record is required.

Coding

The key difference in coding with the XP methodology is communication with the customer. During the coding effort, developers should never try to guess or implement solutions they think meet the requirements of the user story. Instead, they should talk to the customer. This is why XP works. If the developer does not fully understand an implementation, he or she should feel confident in explaining and validating his or her understanding with the customer.

Another key difference is the priority placed on testing. During the coding phase, the developer should always build the test case first, even if the test case is only a written process that can be implemented into a testing tool. Creating the test case first makes it easier to write the code to pass the test case. In addition, because several of the test cases are captured in the user stories, the individual module should be very close to the customer requirements. Another benefit of creating test cases is that it provides an understanding of how the code works. This also means that other developers can ascertain how the code works by reading the test cases.

One of the problems people have with XP is pair programming. This is because it is different. Developers have a

strange ownership issue when it comes to code, and XP tries to break that trend. Remember, the code is written to be rewritten, especially if the refactoring rule is upheld. Therefore, it should be natural for two people to work on the same module. Also, having two people per module will make the system better, because it is built using collaboration. Obviously, only one person can type at a single moment, but more thinking is involved during coding than in actual typing. When developers are having issues with a particular implementation, they have an extra brain to utilize. In addition, two developers are more likely to capture issues and voice concerns to each other, whereas single developers tend to make an educated guess and move forward.

The final component during the coding phase of XP is ensuring no overtime. Several articles have been written about the effects of overtime on a project. Using overtime to complete a project's task tends to tire the team and leave them more prone to making mistakes. Overtime brings down overall moral and provides little if any benefit to the final release. XP provides a mechanism by which to communicate with the customer. Instead of forcing overtime on the team, the release plan should be modified or the scope should be reduced for the current iteration.

Testing

Testing in the XP methodology is one of the main things that sets it apart from other methodologies. Of course, all methodologies require the testing phase. In XP, however, the team thinks about testing from the beginning straight through to the end of the system life cycle. Tests should be automated and included with the actual code. Testing should be completed before the code is integrated and test suites should be created that test all classes and methods within the entire system. The automation of testing is critical because as the

system date approaches, manual tests will be cut short. By implementing testing automation, all tests can be replicated and regression testing can be completed quickly and easily in the maintenance phase of the system. By requiring that all code pass all tests before it is implemented, the development team is ensuring a higher-quality product. Any bugs found during testing should become additional unit tests. Another key difference in testing is that acceptance testing scores should be published. This enables a higher level of quality assurance. The customer will have the ability to review the test results and prioritize the issues found during testing.

Why Use Extreme Programming?


Why should a development team go through the effort of learning yet another development methodology? To answer this question, the benefits of XP should be outlined.

XP is one of the few methodologies that actually strives to deliver high quality code that meets the customer's expectations. Because XP relies heavily on customer interaction throughout the development process, all parties involved take a greater sense of ownership in the final product. This methodology typically converts a group of individuals into a cohesive, high-performance team.

It is amazing to watch an XP team evolve. Since daily meetings are held for 5-10 minutes to outline the goals for the day, all developers are aware of the end goal. These meetings also mean that each development pair is aware of the day's tasks and knows exactly where the project stands. Planning in this way stimulates teams to become more efficient and to perform at a higher level. This strategy is similar to that of sales organizations that put sales goals on a board for all to see. When a team understands where it's going and how it's going to get there it has a higher probability of accomplishing its goal.

“ XP was created to give the customer a method by which to prioritize the requirements for software, thereby enabling the best bang-for-the-buck approach ”

The Last Time You Saw Something This Incredible, It Was Science Fiction



NitroX™ Web Application Development for Eclipse

NitroX's AppXRay™ penetrates all layers of your web application and helps annihilate your web application development problems!

NitroX AppXRay unique features include:

- Debug JSP tags, Java scriplets, jsp: include, etc. directly within the JSP
- Advanced JSP 2.0 and JSTL support
- Advanced JSP editor – simultaneous 2-way source and visual editing with contextual code completion
- Real time consistency checking across all layers (JSP, Struts, and Java)
- Advanced Struts support – source and visual editors for Validation Framework, Tiles and Struts configuration
- AppXnavigator™ – extends Eclipse hyperlink style navigation to JSP and Struts
- AppXaminer™- analyze complex relationships between ALL web artifacts
- Immediate access to variables at all levels of the web application

Download a free, fully functional trial copy at: www.m7.com/d7.do

Copyright © 2004 M7 Corporation. All rights reserved. All M7 product names are trademarks or registered trademarks of M7 Corporation. Java and all Java based marks are trademarks or registered trademarks of Sun Microsystems. IBM, WSAD, WSSD are trademarks or registered trademarks of IBM.



XP is also one of the best methodologies for quality control. This is primarily because testing is put at the forefront of the life cycle. Test scripts are coded before the modules are coded and acceptance testing is required to be error-proof before any code is implemented. XP's strong bias toward testing provides a platform that will allow the development team to build a suite of tests that can be used in the future. This technique removes the fear factor that accompanies code changes in typical organizations. Because the system is designed using tests in this way, new developers can confidently change code within modules without the fear of breaking existing business logic.

A final benefit of XP is customer satisfaction. This methodology enables a customer to see working applications faster than with other techniques. When provided with an evolving system, the customer has the ability to change the prioritization and functionality as business needs change. This enables the software to be implemented in a state where it can provide the biggest benefit to the business. Other methodologies take longer to implement and therefore have an increased potential to miss strategic business opportunities.

Implementing XP into a Development Department

XP implementation is a fairly straightforward task. The three core teams involved are the customer, developer, and testers. Each team plays a vital role in the successful implementation of this methodology. This section will outline the requirements for each team and define a framework to successfully implement this solution. Figure 1 represents the collaboration efforts among the teams.




Figure 1 Collaboration efforts among teams

Customer

The customer plays a vital role in the successful implementation of an XP project. In the initial stages of a project, the customer should be educated on the requirements and time commitments necessary during the system development life cycle. The customer will be required to provide user stories, prioritize tasks, assist with release planning sessions, validate business rules, and assist in acceptance testing validation. This is a large commitment on the customer's part; however, the end product will be much better because of it. One of the biggest benefits of XP, which customers typically appreciate, is the ability to change the product at any time, with minimal cost and little complaint from the development staff. This enables the product to be molded into a system that best fits the business requirements at the particular time of module development.

Developers

The developer plays a somewhat atypical role in an XP organization. The primary difference is that XP does not allow solo programming. Therefore, the "cowboy" programmers who save the day will not be necessary. Developers should be told prior to implementation that the team is going to change into a more collaborative environment. Developers will also need to understand that all code is community code and that they should feel confident in making changes to existing code. The main rules for developers to follow are to keep it simple, collaborate often, work in pairs, refactor judiciously, and work no overtime. Additionally, some developers who are new to iterative programming will need to become accustomed to writing test cases first. This is a culture shift for many developers, but after they try it they will find that the code becomes easier to write and has a higher probability of working the first time. Finally, developers will interact with the customer and the testing team frequently.

Testers

Testers are defined in this example with the understanding that there is an automation department within an organization. Many XP implementations enable the developer to write

all automation scripts for testing, however, this does not take into account other automation tools. I felt it was necessary to break out this group with the understanding that automation scripts will be sent to the testing group, who will use this information to provide statistics for user acceptance testing. The testing group will interact with the developers and the customer on a frequent basis. This team will be responsible for reporting new bugs to the development team and providing information to the customer on issues found during regression testing. All tests should be published. This has a two-prong effect. First, it enables the customer to see and prioritize the issues found during testing. Second, it helps the developer to quickly fix known issues that can be implemented into regression tests.

Framework

1. Start the implementation by educating the core teams on the techniques used in XP. There are many free resources available on the Internet. I have included some links to Web sites at the end of this article. Allow the team a few days to read and comprehend the forthcoming changes. It is also advisable to have a team member who can be the coach for XP. This team member will be the XP expert and will advise the team on any issues that arise. XP is mainly common sense; however, it may take time before some of the techniques become natural.
2. Next, meet with the customer and explain the approach the development team will be taking to implement the proposed software. Customers do not need to understand all of the details of XP, however, they will need to understand the time commitments that will be requested of them.
3. Now that the team has had time to digest the information and the customer understands the requirements, have a kick-off meeting and start writing user stories. The customer will write the user stories with the development team. As outlined in the prior pages, detail specs are not required in initial user stories. Try to keep away from technical information and stay with functional

**IMPORT PEERS.*;
IMPORT EXPERTS.*;
IMPORT YOU.*;**

**// EXPORT DEADLOCK
// GO TO SDN.SAP.COM**



THE BEST-RUN BUSINESSES RUN SAP



You're stuck. You need answers. Maybe you have a solution to share with other SAP developers or a question for an SAP insider. Get the experts, partners and your colleagues to weigh in. Now there's a single collaborative destination where you can all converge: SAP® Developer Network. Nowhere else can you join a spirited discussion forum, download a trial of the latest SAP Web Application Server, take an e-learning course, and, in general, keep us on our toes.

// JOIN IN AT SDN.SAP.COM

“XP is one of the few methodologies that actually strives to deliver high quality code that meets the customer’s expectations”

requirements. Also, remember that user stories should only be about three sentences in length and should require no more than three weeks to implement. If a user story becomes daunting, simply break it into multiple stories.

4. Once all of the user stories have been created, begin to prioritize them. This will be done with the customer. Stories should then be combined into a logical release plan. A release plan is a schedule for the release of given sets of functionality to the customer for validation.
5. Next, have a stand-up meeting with the development team. The development team should meet on a daily basis for no more than 15 minutes to outline their goals for the day. This helps the team understand what the other members will be working on and gives the team an opportunity to request collaboration meetings. Developers will be paired into groups of two. Try to keep the development team at fewer than five pairs. This will make communication easier.
6. Once developer pairs have determined their initial design for a module, CRC sessions should be scheduled. During a CRC session, the developer pair will outline how their particular module will collaborate with other modules within the system. This session is critical, as it allows all team members an opportunity

to understand how the system works. It also provides a platform to locate gaps in design. Developers should keep in mind that simple is better and that fancy functionality can be added later.

7. After the CRC sessions, the developers should write their test cases. These can either be turned into testing code or handed to the automation team that will be responsible for creating the testing automation. Depending on the language being implemented, it may be better to have a process where the test suite is executed prior to migration into the development code base. This will make debugging more centralized. In Java, for example, Ant scripts can be run that will perform all unit tests prior to deployment.
8. Next, the developer pairs will write the code. This is a much easier endeavor with XP because the developers are fully aware of all the test cases and functionality required to implement the given module. They are also aware of all necessary collaboration for the given set of classes they are implementing.
9. Finally, the testing phase should be executed. During this phase, results should be published and bugs should be reported back to the developers, who will in turn create the additional modifications and test cases required to resolve

the issues. The customers will play a vital role in the testing phase, as they will be required to prioritize the issues.

10. After final testing, the code can be integrated into the release. Full user acceptance testing should be completed and the customer should have functionality available for validation. After the release is completed, the customer has the opportunity to modify the prioritization of the next release, as well as change requirements. This cycle repeats until the final product is complete. Figure 2 outlines the XP approach to life cycle.

Conclusion

XP is one of many lightweight development methodologies that attempts to break the barriers and limitations found in the waterfall method of the past. The key difference with XP is that it strives to produce quality code that delivers what the customer and/or business needs. In today’s industries, business requirements change very rapidly. XP attempts to limit the exposure to change by providing a framework designed for change. ☺

References

- *Web sites devoted to explaining the concepts of XP:*
www.xprogramming.com/what_is_xp.htm
www.extremeprogramming.org
- *A Web site that provides a parable for XP:*
www.xprogramming.com/xpmag/kings_dinner.htm

Books:

- Auer, K. and Miller, R. (2001). *Extreme Programming Applied: Playing to Win*. Addison-Wesley.
- Beck, K. (1999). *Extreme Programming Explained: Embrace Change*. Addison-Wesley.
- Fowler, M. et al. (1999). *Refactoring: Improving the Design of Existing Code*. Addison-Wesley.

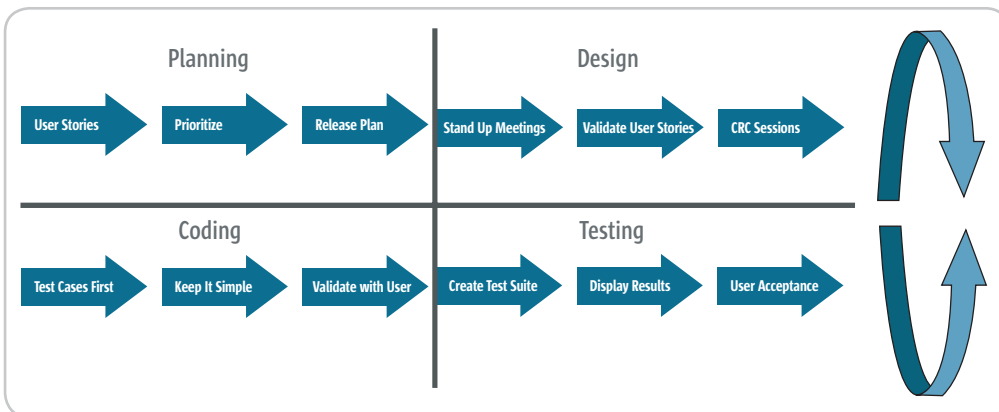


Figure 2 The XP approach to life cycle

The **smart** way to transform XML into Web pages

1. Define XML
2. Map to Web Page
3. Generate XSL and CSS

100% WYSIWYG



XSL IDE

The only comprehensive XSL development platform for efficient making of professional Web pages and Web applications from XML data.
Easy and Intuitive.

XSL Maker

- * Generates all XSL and CSS. No coding required.
- * 100% WYSIWYG editor: build Web page with drag and drop directly in the browser. Position and resize elements to a pixel.
- * Xpath and XSLT wizards: drag and drop code generator
- * Support for secondary XML sources: merge 2+ XML sources
- * Drag-and-Drop Framework for management of XSL variables
- * Rich CSS integrated support
- * Support for Web scripting (JavaScript, VBScript...)
- * Support for run-time XSL input parameters
- * Support for inclusion of other Web pages and Web languages



FREE trial. Download at www.xslmaker.com

Migrating

Enterprise Applications Between J2EE Application Servers

by Ajit Sagar

If you do a search for “migrating J2EE application servers” on the de facto “re” search engine – Google – here are some of the results you’ll get from your query:

- Migrating J2EE applications from Borland JBuilder to IBM WebSphere
- Migrating J2EE applications from WebLogic to WebSphere application server
- Migrating J2EE applications from earlier versions of the application server and from other application server platforms to Sun ONE application servers
- Migrating your J2EE applications to JBoss
- BEA WebLogic to JBoss migration program

When I first got involved in the planning for a project that involved the migration of applications between versions of IBM’s WebSphere product, I naively thought – “This can’t be that big a deal. After all it’s just different versions of the application server. And the changes in the Java platforms between WebSphere versions is well documented. In fact, IBM even provides a Redbook to do the migration...”

I couldn’t have been more mistaken. The migration *was* a big deal. After working on a few of these projects, I can verify that the amount of moving parts in a machine composed of various portfolios, frameworks, third-party vendors, and a variety of stakeholders made the planning for such an initiative a formidable undertaking.

This article covers the aspects of enterprise application migration that involve J2EE application servers, including the motivation, methodology, challenges, and the way to successfully undertake such an initiative. The focus is primarily on the migration of a large portfolio of applications, not individual applications. This article doesn’t get into the basics of application server technologies, Java technologies, etc.; I feel it will be of most interest to architects, team leads, and technical project managers.

What’s the Big Deal About Enterprise Applications?

Many of you have probably migrated applications between versions, hardware and software platforms, etc. The magnitude of the problem increases logarithmically with the number of applications, especially when

they are spread across a number of business portfolios. When we’re considering applications in the J2EE context, enterprise applications typically pose the following challenges:

- The applications are distributed across different tiers – client, business, and database – and they use a different mix of J2EE technologies, including JSPs, servlets, EJBs, JDBC, etc. The versions of the APIs used across the applications are usually not uniform, since they may have been developed at different times.
- The integration requirements of each application may be unique. For example, some applications may require integration with an existing security framework that provides single sign-on. Others may require integration with third-party packaged products.
- Besides the Java platform APIs, the applications probably use other technology components or frameworks that were developed in your organization for Java applications. These may include messaging frameworks, utilities for logging, exception handling, etc. The dependency of applications on such components adds another level of complexity to the migration effort.
- Applications have varying levels of complexity, and it’s hard to apply the same methodology for migration to each of them.
- Application teams have differing levels of expertise, so where the benefits of a structured migration are seen by some, the same is not true for others.

Why Migrate?

Chances are that all your enterprise applications have already been in place for a while and have been working satisfactorily. If such is the case, why migrate? After all, you could continue to support your applications on existing versions of the application server.

Web Services Edge

By the time you get this issue, SYS-CON’s Web Services Edge should be in the office. In fact, that may be where you picked up the copy you are holding in your hands just now. At the conference, we have an App Server shootout event and also a session that I’m hosting on application server migration. If this is of interest to you, be sure to stop by.



Ajit Sagar is a senior technical architect with Infosys Technologies, Ltd., a global consulting and IT services company. He has been working with Java since 1997, and has more than 15 years’ experience in the IT industry. During this tenure, he has been a programmer, lead architect, director of engineering, and product manager for companies from 15 to 25,000 people in size. Ajit has served as *JDJ*’s J2EE editor, was the founding editor of *XML-Journal*, and has been a frequent speaker at SYS-CON’s Web Services Edge series of conferences. He has published more than 75 articles.

ajitsagar@sys-con.com

There are many factors that could have led your organization to consider the migration of applications. One of the main reasons why a migration takes place across the entire organization is because support for the existing version of the application server is being phased out. In such cases, the organization has no choice but to move all the applications to a version that will be supported by the vendor in the future. A typical case is that of IBM WebSphere. IBM plans to phase out version 3.5 of its application server this year. This has led to the mass migration of applications that are housed in 3.5 to the latest version (5.1).

Another reason why organizations end up considering a large scale migration of their J2EE applications is that vendors offer competitive rates/incentives. A typical case is that of Sun offering lucrative pricing on their Sun ONE server to entice organizations into migrating their applications. Another example is that of JBoss. Since this offers the most cost-effective solution, some organizations are considering migrating their J2EE applications to JBoss. For some organizations, it's a matter of consolidating their investment on a standard enterprise platform.

Other reasons include a shift in programming paradigms. Some organizations have recently adopted component-based applications and see a migration to an application server as the natural step to achieve this. Others have realized that departmental applications they have developed cannot scale to an enterprise level. However, in these cases, the effort involved is more along the lines of a rewrite rather than a migration.

I Thought Standardization on J2EE Took Care of Everything

The application server market continues to evolve. Migration is not only about moving non-J2EE applications toward the J2EE standard. Very often, there is a requirement to migrate J2EE applications between application servers. Given J2EE's promise of interoperability, migrating enterprise applications between app servers seems to be more complicated than it should be. This is because application server vendors typically provide extensions to the J2EE platform APIs in order to differentiate their product. In most cases, these extensions were necessary because the performance and rapid development features offered by pure J2EE APIs left much to be desired. For example, when the EJB model was first introduced, it was not usable for enterprise scale applications. It was only in EJB 2.0 that local interfaces were introduced. In the meantime, app server vendors had introduced their own optimizations to address the issues of performance and scalability. When migrating between application servers, the usage of the proprietary extensions adds to the complexity of migration. However, not using the extensions was not the best possible course either, when the applications were originally developed.

Although migration is not only about moving to the J2EE standard, one of the main outcomes of the migration is the move to a standardized platform. This is the case whether the migration is between application server versions, between vendors' products, or between hardware and OS platforms. Figure 1 gives an example of the transition in the case of a migration from earlier versions of IBM WebSphere application server to version 5.1. As shown in the figure, the final state of migration leads to a standardized state of J2EE APIs, IDEs, databases, hardware, etc.

How to Plan for an Enterprise-Level Migration

Enterprise application migration requires careful and detailed planning. In a large organization, at the onset of such an initiative, the planning should be managed in independent tracks that tackle different dimensions of the migration. Figure 2 illustrates the phases and tracks in the planning phase of an enterprise-level migration. As illustrated, these tracks are executed over a number of phases. In the project inception phase, a core migration team should be formed to conduct the migration assessment. As a part of this phase, the team should come up with an initial project plan for the planning phase, determine the methodology and means to communicate with the application teams, identify the key stakeholders, etc. The information gathering phase involves interaction with the various application teams to gather the characteristics of each application. These characteristics include technical parameters such as the number and versions of JSPs/servlets/EJBs, integration requirements, etc., as well as information about releases, schedules, and dependencies between applications and application portfolios.

The analysis phase involves analyzing the gathered information to come up with a feasible migration plan. The final phase is reviews and publication. During this phase, the results of the planning are shared with the stakeholders to then get an agreement on the actual migration of the applications.

Figure 2 illustrates parallel tracks that constitute the migration planning. The architecture planning and strategy track covers the assessment of the architecture, the best possible course for moving toward a standardized J2EE baseline, the migration path to get from the existing architecture for each application to the target architecture, and more. The code migration planning and estimation track deals with the actual migration of the code including the development environment, the steps to migrate from existing versions of the APIs to the target version, the effort estimate for each type of migration (e.g., migration between JSP versions, effort to convert from thick client to thin client, etc.). The test planning and estimation track covers the planning of how to

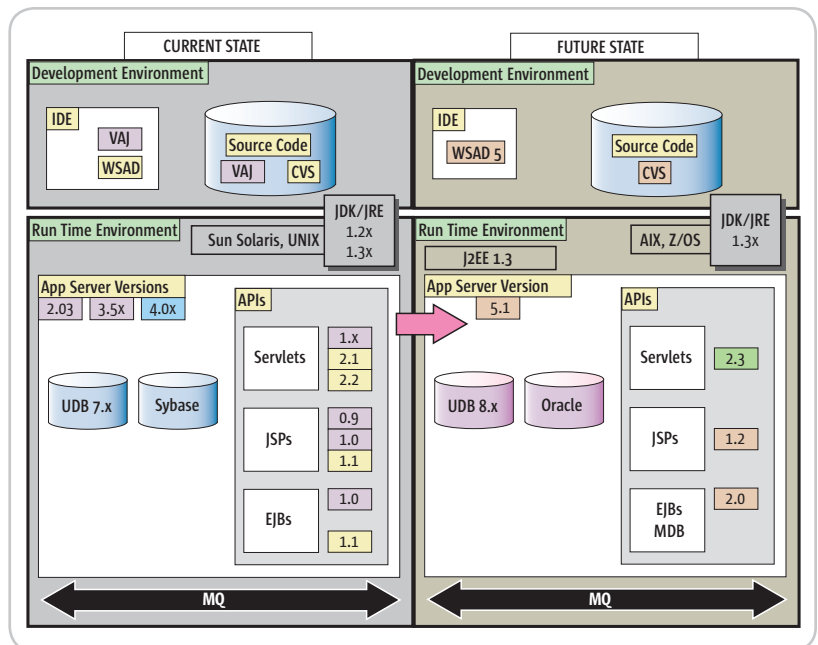


Figure 1 Current and future states of migrating between IBM WebSphere versions

test the application after the migration. The infrastructure planning track deals with the planning of the different run-time environments such as testing, staging, and production. The training plan track addresses the training needs of the development and support teams in order to migrate to the new app server version/vendor.

Best Practices for Enterprise Application Migration

Migrating enterprise applications is a formidable undertaking. However, if planned properly, the migration can proceed smoothly. Some best practices, if followed, can make all the difference between a successful on-time, within-budget migration as opposed to one that can end up being a maintenance nightmare for years to come.

1. Program Manage the Migration

The best possible approach to managing a large-scale migration is to manage it as a separate program rather than treating it as the migration of individual applications. To do this, a core migration team should be created during the planning phase. This migration team should be responsible for coordinating the migration of all the applications. Although the actual migration is best done by the team that owns the application, the core migration team can undertake tasks such as coordination, and training.

2. Prepare Your Arsenal

The good news is that application migration between application servers is a well-defined and understood problem. To that effect, J2EE app server vendors have documented the steps to migrate the code and to deploy the applications on the target environment quite well. In addition, IBM, Sun, BEA, and others have also invested in creating tools that automate some of the effort involved in migrating and deploying the code. This is partly to entice application developers to migrate from a competing vendor and partly to facilitate the phasing out of app server versions that will not be supported in the future. The first step, once you have

identified the need to migrate your applications, is to obtain the sources of information readily available from the vendors and other parties. Read the literature and become familiar with the tools and the quirks of migration that are particular to your product and the J2EE environment platform.

3. Invest in a Planning Phase

If you are migrating a few applications with few inter-dependencies, a planning phase may be overkill. However, in this article the focus is on migrating a large portfolio of applications. To minimize risk and alleviate the impact to the business, make sure you invest in a planning phase. The planning phase should define the characteristics of each application that is being migrated; document the dependencies between applications; estimate the effort in terms of time, money, and resources; produce a roadmap for migration; and produce a step-by-step migration guide for each application that covers all phases of the migration from the migration of code to the deployment of each application in the production environment.

4. Get Everyone Working Toward the Same Goal

An enterprise-level migration involves several parties with different interests and motivations for the migration. A key to success is to get everyone together and working toward the same objective. Ensure that all the concerned parties understand the priorities and rationale behind the migration and understand their role in the initiative.

5. Take the Least Common Denominator

To define a structured path for the migration of multiple enterprise applications, you frequently have to take the least common denominator for defining the move. In other words, all the applications may not see the same benefit in the migration. Some may be in a better position to leverage the target platform's capabilities than others.

6. Invest in Directed Training

When the app server migration is being handled on a large scale, the best way to bring everyone up to speed is to conduct training that is focused purely on the migration. This training can be conducted in-house, or may be provided by the app server vendors themselves. Several training firms have migration tracks that specialize in different types of app server migration.

7. Wade the Waters Before Jumping In

As with all other large-scale projects, your migration roadmap should include a pilot phase. A few applications should be chosen to test the assumptions made in the planning phase such as the code migration effort and performance benchmarks. Once the pilot applications are successfully launched, the plan should be revisited, tuned, and then finally executed.

8. Always Leave a Back Door Open

With ample planning, the migration should proceed smoothly. However, the existing environment should be maintained till the applications are all fully tested and operational. ☺

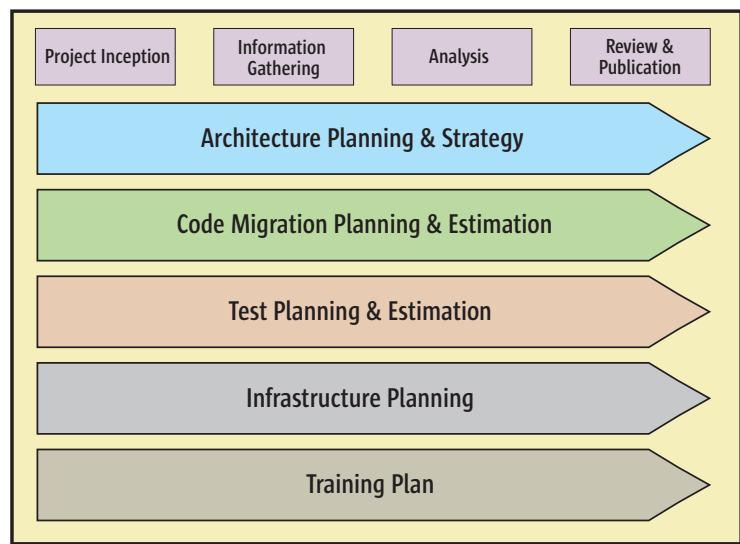


Figure 2 Tracks for planning an enterprise-level migration

Nothing beats our racks

Absolutely nothing



CARRIER CLASS DATA CENTERS

- Highly Secure Guarded Facility
- 24/7/365 Network Operations Center
- 24/7/365 Technical Support
- Redundant Conditioned Air Systems
- Redundant Fiber Entry Points
- Multiple Uninterruptible UPS Systems
- Multiple 1250 KW Generators Onsite
- VESDA Early Warning Smoke Detection

Robert Marsh, Head Surfer

START YOUR OWN WEB HOSTING BUSINESS TODAY!

from
\$299
Instant Activation!

Dedicated Server

Dual Xeon 2.4 GHz
2 GB RAM • 2 x 73 GB SCSI HD
Remote Console • Remote Reboot
2000 GB Monthly Transfer Included

Over 20,000 Servers!

1-800-504-SURF | ev1servers.net

PLESK7
RELOADED
Preferred Control Panel

IP Compliant. Price subject to change. Quantities Limited.
*Per month. Set-Up fees apply. See web site for complete details.



Calvin Austin

Core and Internals Editor

Web Services: Why Can't We Just Talk?

This month's Web Services Edge conference marks four years since the first detailed W3C note on the Web Services Definition Language (WSDL) and nearly five years since the first public specification of SOAP.

You may be wondering, why hasn't the uptake of Web services matched the bold predictions made when it was first launched? There are certainly more developers thinking about Web services with the advent of service-oriented architectures (SOA). However, the number of successful public Web services projects seems to be limited to a few high-profile companies like eBay and Amazon that have published APIs and end points, or developers who have been able to implement services internally but with a narrow and well-defined set of services – still a long way from the smart application that could self-select services at runtime.

There are many factors that can slow the adoption of a new technology. Sometimes it's a technical barrier, competition from an alternative technology, or simply customer reluctance to move to something new. To shine a light on some of the issues that developers experience with early Web services, I can share some of my own experiences with adopting Web services into the J2SE platform.

One common request for J2SE 5.0 was to add Web services, however, even back in 2001 there was some concern over whether Web services was really ready. Unfortunately, we later discovered an issue that would continually dog many other early Web services implementations, namely

interoperability. The issue in this case was that the order of mapping XML back to Java for static stub and dynamic proxy-based JAX-RPC calls was not made a requirement in the specification, and implementers of the specification already had products that didn't match the reference implementation. By its very nature, the Dynamic Invocation Interface (DII) method of calling JAX-RPC was unaffected by this change. The result was that you compiled a Web service with one client runtime, but it may not necessarily work with another JAX-RPC engine. In the case of J2SE, to prevent shipping an incompatible technology, the expert group reluctantly agreed to pull the JAX-RPC specification from the platform.

One of the business problems that Web services was to address was interoperability. Why did the W3C have to create a further Web Services Interoperability profile, and is this going to be the final answer to building a full SOA-based application?

The recent history of computing has been one with specifications drawn up by one or more companies or organizations. Companies and organizations then compete on implementations of those specifications, which then become the resulting standard. Although this has now become accepted practice, this process is not without its flaws.

First, most of the parties that contribute to those specifications have already built or are currently building their own implementations and will often simply modify their own implementation to match the specification. Then, in an effort to be the first to support the new specification, those

same parties will often have to make compromises, especially when it comes to the behavior of the API or other places where the specification is unclear.

To be fair to the specification writers, it's hard just to get an agreement on a well-designed, clean API. What happens inside those APIs, the behavior of the implementation, or what should occur if an API is called in one of the myriad possible sequences is beyond the scope of most specifications.

With that in mind, applying that specification methodology to Web services that are designed for a language- and OS-neutral platform and yet deliver flexibility was never going to be easy. To address some of these pressing issues, the Web Services Interoperability Organization (WS-I) was formed. The first fruit from that collaboration was the Basic Profile 1.0. Last year saw the introduction of an updated Basic Profile 1.1. Yet if you want to use SOAP binding or SOAP with attachments, you need to follow the additional WS-I profiles. The WS-I profiles by their very nature try to define a minimal set of functionality that should work, and as a consequence consign many early Web services technologies, like styles of WSDL encoding, to the trash bin. Even with a restricted set of supported functionality, code examples, and guidelines, the WS-I still makes no guarantees about interoperability. It begs the question: Can Web services fulfill its original promise of dynamic configuration and discovery, or is it yet another useful distributed service that needs to be deployed by architects like the other distributed services that came before it? ☞

A coeditor of *JDJ* since June 2004, Calvin Austin is the J2SE 5.0 Specification Lead at Sun Microsystems. He has been with Java Software since 1996 and is the Specification Lead for JSR-176, which defines the J2SE 5.0 ("Tiger") release contents. calvin.austin@sys-con.com

Outsourcing to Students

—continued from page 6

- Even though students will get a minimum salary for this work, they should also earn academic credits and get graded while working on such projects.

The funny (or sad) part is that the students themselves are already outsourcing their college assignments. There are Web sites where you can hire a coder for any assignment in various programming languages. No job is too small. People from around the world can bid on these projects, and since the offered prices go as low as \$20 USD, it's clear that only programmers from

the developing countries like India or Russia would be interested in these jobs. Academic outsourcing may be even more damaging than the industrial outsourcing, because rich students can improve their grades and earn their degrees without having a good knowledge of the required subjects. Spending more time working as teams in the labs under the supervision of a faculty member or business manager will help minimize academic cheating.

There is one more secret key to the success of commercial projects developed by students: pizza! Each day the client company can send a couple of pies (half plain and half pepperoni) to the labs where the students work. They are going to work for food...and experience. It's a win-win situation for everybody. ☞

Embedded Reporting Options

What You Should Know About Embedding a J2EE Reporting Solution

Development teams have been challenged to deliver applications in record time, in order to meet their organization's ever-changing business needs. This challenge is so great, in fact, that only one-third of these projects are completed on time and within budget.

Over the past seven years, Jinfonet Software has worked with thousands of application developers to understand the challenges they face and the functionality they desire.

The result: JReport 7. JReport is the most sophisticated embedded reporting solution available, specifically designed to reduce the amount of time and resources required to build and maintain reporting functionality in any application. JReport empowers developers to offer their end users a consistent interface for accessing critical business information from within any application.

This article addresses the most common questions asked by developers who need to embed reporting capabilities into their application, and explains why JReport has become the essential tool to increase the odds of successful application development.

How is a reporting solution integrated into a Web Application?

The simplest method to embed a report into a Web application is by using URLs. A tighter and richer integration is achieved by using an API to embed a report into an application's JSPs.

Dynamic reporting functionality can be embedded and customized through the use of DHTML tag libraries. Custom tags help modularize a JSP by encapsulating recurring actions and functions so they can be reused across applications, without custom coding and maintenance.

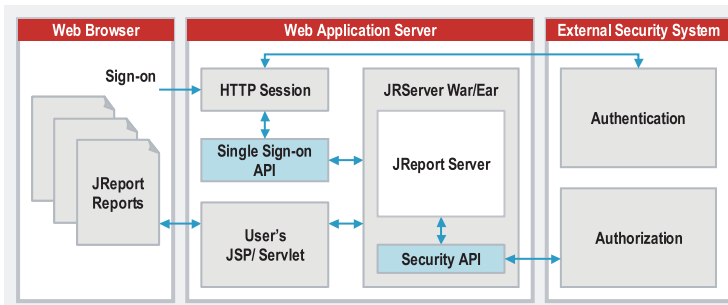
In addition, interactive Web controls can be embedded directly into a JSP or report to provide end users with an easy method for report customization. Web controls can enable any level of functionality, including sorting, filtering, and navigating reports.

How is a reporting solution integrated into a J2EE environment?

A report server should be configurable either as a stand-alone server or integrated with a Web application server as a servlet. The report server can be deployed using various formats including the popular WAR/EAR packages. In some cases, the application architecture may require the report server to be deployed as a separate sub-system. This sub-system can be accessed remotely using RMI calls. An embedded reporting solution

Where can I find the right solution?

As embedded reporting becomes more and more critical, application developers are turning to JReport. JReport is the recognized leader in embedded J2EE reporting. No other company has more years of experience or resources committed exclusively to creating reporting solutions that embed seamlessly into Web applications, provide tight integration with existing security schemes, and offer the rich functionality end users demand from their reports.



should be flexible enough to be configured in different ways, in order to satisfy various requirements.

How is security integrated?

An embedded reporting solution should provide a security system that controls data access at the record-, column-, group-, page- and report-level. The reporting solution should also provide flexible ways to integrate directly with existing security schemes. A popular method is to invoke external security systems like LDAP. This links the reporting solution into existing security systems.

How do the integrated application and the reporting solution provide single sign-on capabilities?

Ideally, an application and its embedded reporting solution will work as a single, seamlessly integrated system. Usually, it is desirable for the reporting solution to connect to an external security application for authentication and authorization. This method allows for login at the application server level, providing the system with single sign-on capabilities.

What are developers saying about JReport?

"We got to market fast, and our customers found JReport to be an easy-to-use, powerful solution for developing applications with strong reporting capabilities"

- Unify

"JReport is an effective, highly scalable, and easy-to-embed solution that has enabled our application to provide significant value and cost savings."

- Storeroom Solutions

"JReport was the best embedded solution for us. JReport integrated easier, scaled better, and was easier to customize than any other solution."

- Bell Canada

To learn more about how JReport helped these companies, and for a FREE trial, call 301-838-5560 or visit www.jinfonet.com/jp2.htm.



JReport is a 100% Java reporting solution that leverages J2EE industry standards and modular components to deliver a solution that easily embeds into any application. With JReport, report development and deployment in diverse environments no longer require costly integration of various resources from different vendors, various data sources, and legacy systems.

Jinfonet Software
(301) 838-5560
www.jinfonet.com

Opening Java Projects up to Open Source

by Timothy Fisher

The open source community has room enough for all

As open source technology is gaining more popularity in the press and among the general population, there still seems to be a lack of knowledge of what is available via open source amid many software development projects. While the mainstream media and the average computer user thinks of Linux whenever the term open source is brought up, this article looks at the amazing wealth of technology available to Java software development teams. For years there have been promises of code reuse resulting from the use of object-oriented programming. The open source community is delivering on that promise today.

If you're a Java developer, there is a gold mine of good open source technologies available now to assist you with your projects. There is an open source project covering nearly every category of technology, component, or framework that you can envisage.



Timothy Fisher has

recognized expertise in the areas of Java, XML, and security technologies. He has served in technical leadership and senior architecture roles with companies such as Motorola, Cyclone Commerce, and PricewaterhouseCoopers. Currently, he is employed as a senior Java technologist with LogicaCMG in Southfield, Michigan. Timothy is also a frequent writer and speaker on Java and security technologies.

trfishermi@yahoo.com

The Communities and Projects

Apache Software Foundation

The Apache Software Foundation, best known for its Apache Web server, is probably the most well-known source for open source projects. Apache products are widely considered to be best-of-breed in their categories. Apache's Jakarta project is home to some of the best Java development technologies available. In the Java community, one of the most well-known open source projects is the Struts project, which is part of this Jakarta family. It can be arguably stated that Struts has become a de facto standard framework for implementing Java Web applications.

One of the few open source development projects in wider use is the Ant build system, which is another project from the Apache team. Ant is a robust, enterprise-ready tool for building and deploying your Java applications. Any Java project that is not using Ant as a

part of its build process today should be doing so. Apache is also the home of the most widely used Java XML libraries, Xerces and Xalen. It's hard to find a Java project using XML that doesn't make use of these libraries.

Another popular open source project from the Apache group is the Tomcat server. Tomcat is the official reference implementation of the Java servlet and JSP specifications. Tomcat is a capable server and has widespread usage, particularly as a development test server. The most significant piece of J2EE not supported by Tomcat is the EJB specification. Many projects that don't use EJB technology will use Tomcat as a testing environment on local developer workstations, even if the production server will ultimately be a commercial J2EE product such as WebSphere or WebLogic. Tomcat is also used in many production environments.

JBoss, Inc.

If you want a full J2EE application server that does give you access to EJB technology, the open source community has a powerful offering that meets your needs as well. The JBoss server, created and supported by the JBoss group, is a standards-compliant J2EE application server. The JBoss group is proving that open source can also be adapted to a successful commercial business model with its recent announcement of 10 million dollars of venture funding led by the venture capital firm, Matrix Partners.

The Apache group has a competitor in the works as well. The Apache Geronimo project will be the J2EE application server from the Apache group. Many of the Geronimo developers are former JBoss developers. As of this writing, Geronimo had not yet released a J2EE-compatible server, but rather what they call "technology demonstrations." Geronimo is a great example of what will be an enterprise application built on


top of many other open source projects. Rather than build from the ground up, the Geronimo team is wisely making use of many open source projects already in existence. This allows the Geronimo team to rapidly build an enterprise-class application server in a time frame that would be daunting to any commercial development team.

JUnit

One of the hottest trends in application development today is the concept of test-driven development. This philosophy places a new emphasis on creating unit tests for all of your methods as you write them. Here's another area where open source shines. The de facto standard unit-testing framework for Java development is JUnit. JUnit greatly simplifies the creation and execution of unit-level tests for your code. As I've stated about Ant, if your project is not using JUnit, it should be.

SourceForge

While Apache is probably the most well-known place for open source Java technology, the biggest resource for developers is the SourceForge community Web site. SourceForge is the home of thousands of open source projects. At the time of this writing, they had over 12,000 Java projects listed on their site. Not that these are all active and useful projects. Many of them are lying stagnant with outdated code, or no files have been produced. Because of this, when you do a search on the SourceForge site, don't get too excited about all the results until you've had the chance to actually browse the details of the projects to make sure they contain useful technology. For those developing open source projects, SourceForge provides free hosting services for your projects. SourceForge is usually the first place to check if you are looking for an open source component covering a particular technology.



:: Application Server Configuration ::

[Apache WebServer Configuration](#)
[App Server Virtual Hosting](#)
[Heartbeat Load-Balancing & Clustering Help?](#)

AppServer Type: JAVA J2EE JAVA J2EE
Present AppServer: Tomcat 5.5.4
AppServer Virtual Path(s)/Extension(s): /*.jsp /servlet/ /manager/ /*.do

Select Application Server: Tomcat 5.5.4

Customize Application Server:

- Tomcat
- Tomcat
- JBoss-Tomcat
- ColdFusion-MX-Enterprise
- Resin
- JRun
- JBoss-Jetty
- Jetty
- JBonAS-Tomcat
- CUSTOM
- NONE

Localhost Bind Port Range: 4671-4679
AJP/HTTP Bind Address: localhost:4671
Dedicated IP/Host Address: myaccount.webappcabaret.net
App Server Home Directory: /usr/ngasi/contexts/myaccount/e
Java Virtual Machine (JVM): JDK1.5.0
Connector: AJP

Imagine a hosting company dedicated to meet the requirements for complex web sites and applications such as those developed with Java J2EE. At WebAppCabaret our standards based process and tools make deploying Java J2EE applications as easy as a point-and-click. We call it **Point-and-Deploy Hosting**. Our advanced NGASI Web Hosting management Control was designed for the hosting and management of complex web sites and applications thus cutting down on maintenance time.

Backed by an experienced staff as well as a **Tier 1 Data center** and network. Our network is certified with frequent security audits by reputable security firms.

All hosting plans come with **advanced tools** to manage your application server and Apache web server. Not to mention the other features, such as virus-protected email, bug tracking, and many more portal components.

Complete power and control at the tip of your fingers. We take care of the system and hosting infrastructure so you can concentrate on development and deployment of your application. That is **real ROI**.

Log on now at <http://www.webappcabaret.com/jdj.jsp> or call today at 1.866.256.7973



The SourceForge site is divided into subcommunities called foundries. Each foundry is a community focused on a specific technology. Our primary interest here is in the Java foundry. The Java foundry is a central location where you can find news, announcements, discussions, and articles about the Java projects hosted by SourceForge.

Tigris.org

A lesser-known site that provides good open source projects is the Tigris.org site. Tigris.org is an open source community focused on building better tools for collaborative software development. While not nearly as large as the SourceForge community, a nice thing about the Tigris site is that all of its projects are active and ongoing. Thus, you won't come across any dead or barely started projects with no usable code as you often do on SourceForge. The Tigris projects don't have a complete Java focus, but rather tend to be focused on more of the noncoding-related tasks that are part of

news, and blogs related to Java technology. This is a great educational site for those new to the Java platform and those who wish to expand their knowledge of the platform and technologies available. With all of its offerings, this is a well-rounded collaborative community site where projects, discussions, and other offerings are divided into communities of interest related to specific technologies under the Java umbrella. This is a nice distinction that this site offers. If you're looking for a project in a specific technology, say a Java networking project, you can search within that specific community. For an experienced Java developer, a periodic review of this site will help you keep current on emerging Java technologies.

Open Symphony

Open source J2EE enterprise class components can be found at the Open Symphony project. This is a project that emphasizes quality over quantity. You won't find nearly as many components and technologies available here as on

Codehaus

Another open source community site hosting Java projects is the Codehaus site. Codehaus is an open source project repository with a strong emphasis on Java projects. On its home page, Codehaus specifically states that all of its projects are business friendly in terms of licensing. Codehaus also has some major corporate sponsors, including BEA. Codehaus focuses purely on providing projects. You won't find the community features here that you find on similar sites. One of the most popular projects on Codehaus is the OpenEJB container. This is an open source implementation of the EJB specification. The Apache Geronimo project has announced it will use the OpenEJB container as its EJB container, giving this project significant visibility as well as a chance to gain Sun certification. Other popular projects hosted by Codehaus include AspectWerkz, an AOP framework for Java; Drools, a business rules engine; and PicoContainer, a lightweight component container.

“ If you're a Java developer, there is a gold mine of good open source technology available now to assist you with your projects ”

the software development cycle.

A particularly good project from Tigris is ArgoUML. ArgoUML is a full-featured UML modeling tool implemented in Java. For those who are not familiar with ArgoUML and are in need of a good software modeling tool, check out ArgoUML. I think you'll be impressed with what you can get for free. One of the most intriguing and perhaps most useful projects that Tigris offers is a library of reusable software engineering document templates in the ReadySET project. Here you'll find templates covering all your software project deliverable documents, including project planning, requirements, architecture and design, and deployment and test documents.

Java.net

New this year to the open source community space is Sun's Java.net community. This is really much more than an open source community. The open source projects are just a piece of the Java.net community. This community also provides articles, discussion forums, mailing lists,

some of the other open source communities, but those you do find are very robust, scalable, enterprise-ready components that have been used in many commercial applications. In their own words, their components stress simplicity, integration, plugability, and specification compliance. Popular components available here include a good workflow engine, OSWorkflow, and WebWork, an MVC framework alternative to Struts. The components here are all developed by the same core team, and thus they cover a spectrum of enterprise technologies. Unlike some of the larger community sites, you won't find 10 components here trying to accomplish the same thing. For the most part, these components are designed to complement each other and give the developer the building blocks necessary to develop an enterprise-scalable application. There is a very unrestrictive common license adopted for all of the Open Symphony projects as well, allowing for source and binary redistribution and/or modification in both commercial and open source projects.

Java-source.net

The sites mentioned so far all provide hosting of open source Java projects. A very good site with a slightly different objective is the Java-Source.net site. Unlike Apache, SourceForge, and Tigris, Java-Source.net doesn't host or sponsor any projects. Instead, this site serves as a directory of open source Java technology, components, and applications. This site is an excellent starting point when you're searching for Java components in a specific category or just browsing the available projects.

FreshMeat.net

Another directory-like site to be aware of is the FreshMeat.net site. This site provides a directory of Unix and cross-platform software. It claims to have the Web's largest index of this software. While it does list some commercial projects, the directory is dominated by open source projects. Unlike the Java-Source.net directory site, here you can also find community features such as user-contributed articles and discussion boards. FreshMeat is owned

by the Open Source Technology Group (OSTG Inc.), which is the same company that owns and runs SourceForge.

onjava.com

The O'Reilly publishing company has long been a strong advocate of open source technology. Its onjava.com site features excellent open source content with lots of articles on popular Java open source projects. This is a good site to go to for education.

Eclipse

With the rise of the Eclipse development platform created by IBM, even the IDE domain, once completely dominated by commercial products, is being overrun by quality open source technology. The most popular use of the Eclipse platform is as a powerful Java IDE. The Eclipse Java IDE shares much of the same codebase as the commercial IBM WebSphere Studio Application Developer IDE product. In addition to achieving high ratings from developers, Eclipse is rapidly grabbing market share away from commercial competition.

We've mentioned some of the most popular open source community sites, but there are also projects that are not hosted or sponsored by a community site but have either their own sites or hosting on a corporate site. The best way of finding these projects is through a Google search. Usually if you type in the category of project you are looking for along with the term *open source*, you'll get plenty of hits for open source projects. Many of these projects are also listed on sites that try to be directories of open source projects, such as JavaSource.net and FreshMeat.net.

Persistence Frameworks

One of the biggest and fastest-growing areas of open source technology in the Java space is in the area of persistence frameworks. Some of the top contenders here include Hibernate, Cayenne from the ObjectStyle Group, Apache's Torque, and Castor from the Exolab Group. If your application contains anything beyond a small amount of database code, these frameworks can be very beneficial to you. As we all know, writing pure

JDBC and managing the connections, transactions, and exceptions can eat up a lot of development time. These frameworks assist you in all of those areas. If you have a large database with a large number of tables, it can be burdensome to generate data access objects and the code necessary to read, update, create, and delete data rows and relate them to your data access objects. An open source persistence framework such as Cayenne will automatically generate all of your data access objects, including all of the code necessary to create, read, update, and delete objects directly from your database schema. This feature alone can save hours of repetitive, error-prone programming.

Breaking Down the Barriers to Open Source

Occasionally I'll come across a project team whose members will say they don't use any open source technology because their management is opposed to it, or there is some corporate guideline against its use. This is more common in high-risk industries such as banking and

Think grid computing is only for specialized applications? **Think again!**

xTier™ / GRID
Grid Computing for Java - Made Simple

Bring the power of GRID
to your Java/J2EE applications

Fitech Laboratories Inc.
sales: +1.646.495.5076 or sales@fitechlabs.com
support: support@fitechlabs.com - 24/7/365

www.fitechlabs.com/grid

finance. These types of broad anti-open source guidelines are created by managers or technologists who are uneducated with regard to the real benefits and risks of open source, particularly in terms of open source technology for developers. Those who create restrictive open source policies need to stop equating the issues of the usage of open source end-user applications with open source development components and frameworks. Unlike the case of end-user applications, when comparing open source developer technologies to commercially available developer technologies, in most cases with open source you have more, not less, support available; greater, not less, ability to establish trust in the components; and greater, not less, ability to get timely turn-around on bugs and feature requests. Any corporation that has a policy concerning open source that does not distinguish between end-user applications and developer technologies will end up with these overly restrictive policies that are strategically and financially harmful to the corporation.

When a company adopts overly restrictive policies regarding the use of open source technology components, the reasons for adopting those policies are typically variations of the following:

1. We don't trust the source code, i.e., it could be full of bugs or security holes.
2. There is no commercial support for it and, if we encounter a problem with it, we have no one to turn to.
3. We have a legal policy prohibiting us from using open source due to license issues with open source.

Let's address each of these reasons. The available technical support on many of the open source projects is far superior to the support you get with many commercial products. With open source projects, typically you're able to directly contact the developers, and many of them also have very active message boards and mailing lists that provide great technical support. Even if you assumed that absolutely no support is available (though I haven't found a useful open source project for which this is true), I'd still argue that having your team learn and understand the source code of the open source project well enough to support it will still save you a great deal of time compared to rewriting the equivalent technology yourself.

There is also an increasing number of consulting companies offering support for open source projects.

Regarding the issue of trusting the source code of open source projects, don't overlook the fact that an advantage of open source projects is that you have access to the source code, and if there are areas that you don't trust, you can review those areas to gain whatever trust you feel is necessary. The same can rarely be said about commercial, closed source projects.

The fact that open source projects, especially the more popular ones, have been subjected to review by many development teams and individuals using them backs up the argument that these projects are often more secure, more bug-free, and thus more trustworthy than their commercial counterparts.

In the case of a commercial project, you must place your trust not in the source code, which cannot be verified by your team, but in the commercial vendor of the product. While on a commercial project the source code may have been reviewed by a relatively small team of developers, an open source project is open to review by thousands of very capable developers. If there is a particularly sensitive piece of functionality that you want to review more closely, with an open source project you or your team can thoroughly review the source code line by line. With a commercial product, unless they are willing to show you that piece of source code, you again must rely on the reputation of the vendor and not the code itself. You'll often hear that relying on the reputation of the vendor is an adequate assurance of trust because, unlike with open source, the vendor may be financially liable for problems in its code. My first response to this is to ask yourself this question: Are you sure that the vendor is liable? In most cases the licenses of the commercial products specifically remove liability from the vendor for bugs and security flaws. Also, to counter this argument, we have only to look at Microsoft. Here is a company that, based on the commercial liability argument, should have been producing very bug-free and secure code. Yet we all know that Microsoft products have become notorious for their security flaws. Even Microsoft is finally beginning to understand that opening its source code, as it is doing with parts of Windows, may

actually help make their products more secure and more trustworthy.

The security world is way ahead of the rest of us. In the realm of security, it has been a long-standing belief that cryptographic algorithms and security protocols that are published and open to inspection are much more secure and trustworthy than proprietary algorithms and protocols. We need to start believing that with regard to nonsecurity software components as well. Openness is a good thing.

Many large corporations will allow a small development team, often composed of temporary contractors, to develop core functionality, and they will trust that more than they are willing to trust an open source project. This occurs despite the fact that the open source project in almost all cases has had much more time to mature, and is often reviewed by thousands of developers and development teams. This is simply an illogical policy based on uneducated decision making. These companies spend countless hours developing functionality that is rarely as good as what could be obtained in the open source community.

The Linux lawsuits brought by SCO and the media attention surrounding them are instilling fear in many lawyers, causing them to issue blanket policies regarding the use of open source in their organizations. Based on legal department opinions, some organizations are dictating policies that disallow developers from using any open source projects because of potential legal problems and license issues. This is equivalent to having your legal department say that you can't use any Microsoft software because of security problems with Internet Explorer. The legal departments making these policies don't really understand open source and the variety of licenses commonly used in open source projects. Rather than take the time to understand the licenses and treat these projects independently, they issue these restrictive policies that are ultimately bad for the business.

Another way open source projects can benefit your development team in environments that face these restrictive open source policies is to not necessarily adopt complete open source projects, but to reuse parts of the open source projects. You'll find the licensing of

many of the projects is very amenable to this type of use. With this technique, you can make the argument that you're not "using an open source technology," but rather you're borrowing the code from some of these projects and making it part of your own components, which you fully support and maintain. There are many open source projects that are very valuable even if you don't use them in their entirety, but rather in a piecemeal fashion. Pick and choose the pieces you need.

Wrap Up

In addition to using open source Java technology on your professional projects, open source projects are an excellent educational resource. If you are trying to learn a new technology, there is most likely an open source project available for download that you can get familiar with and thus learn by example. Once you think you understand the project, you might even see if you can modify it and create a useful piece of functionality that can be contributed. Often, this is how open source continues to grow and flourish.

The use of some of the open source frameworks is the area in which you'll see open source technology have the greatest overall impact on your project. The widely used frameworks available in the open source community have gone through many iterations of improvements, design reviews, and bug fixes. Frameworks such as Spring, Struts, Cayenne, and Hibernate are well-architected projects that can help you improve the overall architecture and design of your project. If I have a large J2EE enterprise application, and I'm able to tell developers joining the team that we use Struts as an MVC framework, Spring for middle tier management, Hibernate as a persistence framework, JUnit for unit testing, and log4j for logging, I've given that person a fairly powerful overview of our technical architecture. It also happens that there is at least one published book available on each of these components, so, even aside from any project documentation, there is a foundation for learning and understanding your project.

Before choosing any open source code for inclusion in your own project, be sure that you have read and understood the terms of the license. Although many of the open source projects, including most of the more popular ones, have a very relaxed license that essentially allows you to make use of the code in any way you see fit (be it educational or commercial use), there are some open source licenses that require any derivative products to also be made open source. This type of license is, of course, not compatible with a project you may be developing for commercial sale, and thus these are the open source projects to avoid.

Before you begin your next project, it's in the best interests of your team and organization that you take the time to investigate open source projects that may be of use to your project. You'll be pleasantly surprised by what you find. As software developers, we need to start realizing the promise of code reuse and stop reinventing some pretty expensive wheels. ☺

Communities and Directories

- *Apache Jakarta*: jakarta.apache.org
- *Codehaus*: www.codehaus.org
- *Exolab*: www.exolab.org
- *FreshMeat.net*: www.freshmeat.net
- *Java.Net Community*: www.java.net
- *Java-Source.net*: www.java-source.net
- *Object Style*: www.objectstyle.org
- *Open Symphony*: www.opensymphony.com

- *O'Reilly*: www.oreilly.com
- *SourceForge*: www.sourceforge.net
- *Tigris.org*: www.tigris.org

Projects Mentioned

- *Ant*: ant.apache.org
- *ArgoUML*: argouml.tigris.org
- *AspectWerkz*: aspectwerkz.codehaus.org
- *Castor*: www.castor.org
- *Cayenne*: www.objectstyle.org/cayenne
- *Drools*: drools.codehaus.org
- *Eclipse*: www.eclipse.org
- *Geronimo*: geronimo.apache.org
- *Hibernate*: www.hibernate.org
- *Jboss*: www.jboss.org
- *JUnit*: www.junit.org
- *Log4J*: logging.apache.org/log4j
- *OpenEJB*: www.openejb.org
- *OSWorkflow*: www.opensymphony.com/osworkflow
- *PicoContainer*: picocontainer.codehaus.org
- *ReadySET*: readysset.tigris.org
- *Struts*: struts.apache.org
- *Tomcat*: jakarta.apache.org/tomcat
- *Torque*: db.apache.org/torque
- *WebWork*: www.opensymphony.com/webwork
- *Xalen*: xml.apache.org/xalan-j
- *Xerces*: xml.apache.org/xerces2-j

COMMON CONTROLS www.common-controls.com

The Java™ Presentation framework for J2EE™ Web applications

Based on:

- Java™
- Servlets™
- Java Serverpages™
- and Struts

For Free! Get your free trial version – www.common-controls.com
See the common controls in action – go for the Online Demo!

Contains the most common control elements which are required for the development of J2EE™ applications with rich HTML frontends like:

Lists Trees Tabfolders Menu Forms
BreadCrumbs Calendar Colorpicker

www.common-controls.com

Web Services and XML Special

*A Java Specification Request Update from the
Java Community Process Program*

by Onno Kluyt

*If you're reading this at the Web Services Edge 2005
East Conference, you're probably having a great time.
If, on the other hand, it's before February 15th, you might
still have time to get to Boston and attend. For anyone
interested in Web services or XML, this conference will
be big.*

Whether you're at the conference or not, I want to highlight several of the newest JCP Program Java Specification Requests (JSRs) that promise to improve Web services development and deployment. Of the 20 or so JSRs that offer benefits for Web services, I've described what's happening on several that might be of interest to you. I apologize in advance if I missed your favorite.

Steady As It Goes

The Java platform will celebrate its 10th birthday later this year, and it has evolved as an undeniable foundation for secure, compatible Web and networked services. In fact, I believe the Java platform is the most secure, reliable, and supported architecture for Web services and a service-oriented architecture (SOA).

The JCP program joint community effort is the key to what's been achieved in maintaining the platform's compatibility, interoperability, and flexibility for Java technology and Web service developers. There are challenges, obstacles, and hurdles ahead, but it's great fun to take a look at the work being accomplished by the JCP program membership.

Let's check out a dozen or so perfect examples now.

JSR 173

The Streaming API for XML is a new method for pull-parsing XML; it allows the programmer to ask for (or pull) the next event. "It's bidirectional, supporting reading and writing, and is designed to integrate with JAX-RPC," says Chris Fry, spec lead from BEA Systems. "Parsing XML for high-performance streaming and the associated API grew out of BEA's XML and Web services team." The reference implementation is open source, and is available at codehaus.com: <http://stax.codehaus.org/>. There's also a mailing list for anyone interested, StAX Builders at Yahoo! Groups: http://groups.yahoo.com/group/stax_builders/.

JSR 181

Web Services Metadata for the Java Platform makes Web services development easier to learn and do. "It's a spec for creating Web services using the new JSR 175 annotations, and it's setting a precedent among other JSRs for using annotations in a declarative programming model," says Brian Zotter, spec lead from BEA Systems. "We really think this is going to be the development model of the future." JSR 181 is in the final stages of the spec headed for final approval in a matter of months. Catch the JSR 181 messages at Yahoo! Groups: <http://groups.yahoo.com/group/jsr181/>.

JSR 262

Web Services Connector for Java Management Extensions (JMX) Agents is now open and inviting nominations for Expert Group membership at <http://jcp.org/en/jsr/egnom?id=262>. Eamonn McManus, spec lead, says this JSR will define a connector for the JMX Remote API that uses Web services to make JMX instrumentation available remotely. Clients don't have to be Java applications, but they can be.

“The Java platform will celebrate its 10th birthday later this year, and it has evolved as an undeniable foundation for secure, compatible Web and networked service”



Onno Kluyt is the director of the JCP Program Management Office, Sun Microsystems.

onno@jcp.org

JSR 208

Java Business Integration (JBI) extends J2EE and J2SE with service provider interfaces (SPIs) for business integration. These SPIs will enable the creation of a Java business integration environment for specifications such as BPEL4WS and the W3C Choreography Working Group.

The actual SPIs will probably be of great interest to integration technology implementers, even more so than developers. Developers will be more interested in the ability to plug best-of-breed components, often based on disparate technologies, to create a service-based architecture. This standard enables a composite view of an SOA-based enterprise for the first time. In addition to the SPIs, a JBI-compliant platform will provide crucial enterprise services such as a normalized message service, protocol bindings, and service engines. Recent additions to the Expert Group include Apache Software Foundation, Fujitsu, and JBoss.

JSR 244

The J2EE 5.0 Specification is expected to be finalized in Q1 of 2006, and enterprise developers are watching this release very closely. It's focused on J2EE ease of development. It will embrace the needs of both sophisticated and beginning enterprise developers, and, of course, support the latest Web service standards. If you're doing any kind of enterprise development for Web services, J2EE 5.0 is definitely a release to explore.

Two of the most exciting additions to J2EE are JSR 181 and EJB 3.0. JSR 181 provides a much easier way to create Web services using simple annotations on Java code. EJB 3.0 greatly simplifies the EJB programming model and introduces a new Java persistence API that makes it much easier to map Java classes to information in a database.

JSR 267

Take a look at the *JavaServer Pages (JSP) Tag Library for Web Services*. A standard JSP tag library for accessing Web services would be quite helpful, and this should be finished for final release near the end 2005. It will help JSP authors integrate Web services into Web applications with much less effort and quicker turnaround times. The benefits are strong, and it's a compelling assignment, well worth the effort, supported by JBoss, Oracle, and Sun Microsystems. Stay tuned for the public review scheduled for July 2005.

JSR 206

Java API for XML Processing (JAXP) 1.3 is a dot-release update to JSR 63, JAXP 1.2, both of which enable applications to parse, validate, and transform XML documents. This version accommodates the newest revisions to several core XML standards, including XML 1.1, Namespaces in XML 1.1, Document Object Model (DOM) Level 3, and Simple API for XML (SAX) 2.0.1. It's included in J2SE 5.0. The JAXP 1.3 design centers are standards compliance, performance, and ease of development. It introduces an object model agnostic XPath 1.0 API and a validation API that supports the reuse and caching of schemas.

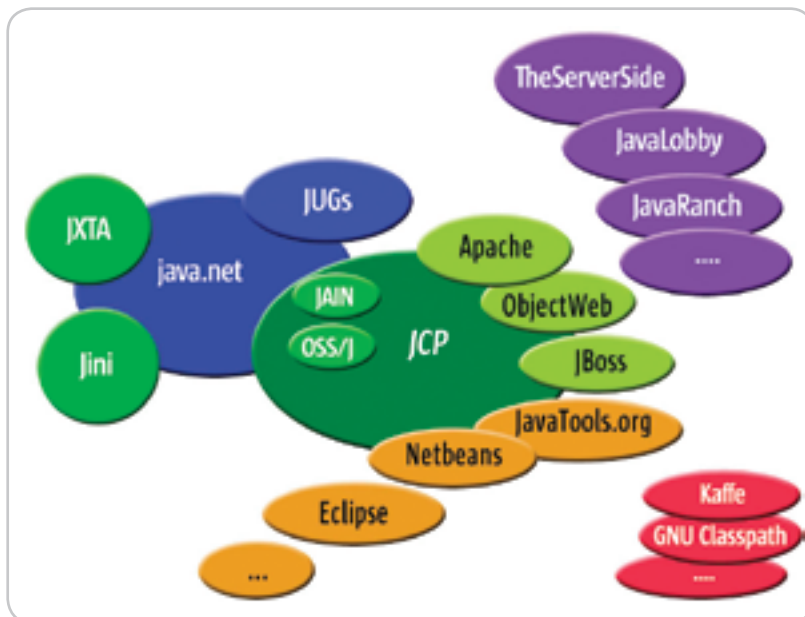
JSR 222

Java Architecture for XML Binding (JAXB) 2.0 is a follow-up to JSR 31, JAXB 1.0. It simplifies development by binding XML documents to Java objects based on a document's

While You're at Web Services Edge East 2005 Conference

At the Hynes Convention Center, Boston, February 15-17, 2005

- **JCP Training & Communication Session**, Monday, February 14, 9:00 a.m. – 6:00 p.m. (Pre-conference) led by JSR program manager, Liz Kiener
- **Developing Web Services Using Java Technologies**, Tuesday, February 15, 8:00 a.m. – 11:00 a.m. Attend the free JCP Tutorial led by JCP PMO manager, Aaron Williams. This tutorial will discuss many of the JSRs featured in this article and will fill you in on how to become more involved with these technologies.
- **Web Services and Security – Java Technology Panel Session**, Tuesday, February 15, 3:00 p.m. – 3:50 p.m. I will lead this panel discussion comprised of several JCP Program Executive Committee Members.
- **The JCP Program: How the Java Technology Binary Software Standard Is Managed and Evolves**, Thursday, February 17, 4:00 p.m. – 4:50 p.m. I will also be presenting this Java track session.
- **Visit the JCP program in the Expo Hall**, Booth #405
Meet with members of the JCP Program Office on Tuesday, Wednesday, and Thursday afternoons during the conference.



The JCP program is part of the overall Java ecosystem, responsible for maintaining the standards and compatibility of the Java platform. Those core promises – standards and compatibility – help grow the developer community and the marketplace.



JCP program members and the public are encouraged to review and comment on the Expert Group's specification draft during the Java Specification Request (JSR) process.

The average JSR is finalized for release after collecting JCP program member and public comments during the Early Draft Review and the Public Review. Every step is consistently documented and available to members.

The JCP Program Today

JCP program members collaborate to define Java technology as a trusted binary standard and the foundation for compatible implementations.

The current version of the JCP program, JCP 2.6, initiated in March 2004, greatly expanded community participation in an open, inclusive process:

- Encourages broad individual, corporate, academic, institutional, government, and non-profit participation
- Gives spec leads their choice of technology licenses, including open source.
- Guarantees public specification availability and feedback during development.
- Offers all specification drafts to anyone with an Internet connection for review and comment.

JCP Facts at a Glance

- 850 members
- 267 JSRs, about one-third of which are complete
- Members initiate about 40 new JSRs a year
- JSRs are led by over 55 JCP member entities

From the JCP...

Web Services: Java Web Services, XML, 20+ JSRs

Enterprise: Two versions of the J2EE platform

Desktop: Two versions of the J2SE platform

Wireless: Two versions of the J2ME environment – CLDC and MIDP

Vertical Industries: OSS and JAIN APIs

Tools Community JSRs, IDE compatibility

schema. The binding enables the manipulation of XML content at the object rather than at the parser level. JAXB 2.0 has significant new functionality, including support for all XML schema, the binding of existing classes to XML schema, and data binding support for JAX-RPC 2.0 for Web services development.

JSR 224

Java API for XML-Based RPC (JAX-RPC 2.0) is, in short, an easy-to-develop programming model that helps you build Web service clients and endpoints based on SOAP. JAX-RPC clients can invoke Web services developed across heterogeneous platforms. Similarly, JAX-RPC Web service endpoints can be invoked by heterogeneous clients. This update supports a lot of new functionality: JAXB 2.0-based data binding, SOAP 1.2, and standardized metadata for Java-WSDL mapping. You'll also notice its new ease-of-development features, support for easier evolution of Web services, an improved framework, and support for asynchronous RPC and non-HTTP transports.

JSR 265

The API for Utilizing Web Services Policy is led by SAP. This specification aims to standardize an API for utilizing Web services constraints and capabilities. Constraints and capabilities refer to Web services metadata, such as quality of service (QoS). It will provide interfaces to obtain and access the capabilities in effect for any given message exchange instance. This API will enable other Java technologies to utilize constraints and capabilities in a uniform manner, thereby strengthening Java as a competitive Web services platform.

It's expected that JSR 265 will be utilized from within other Java specifications for Web services, such as JAX-WSA (JSR 261), JAX-RPC (JSR 224), and JBI (JSR 208).

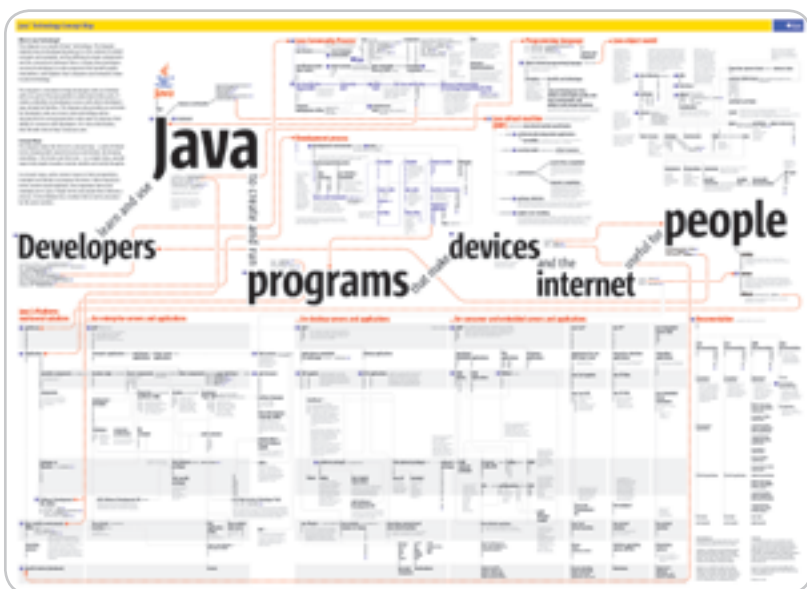
JSR 109

Implementing Enterprise Web Services is an IBM-led spec that's part of the umbrella J2EE 1.4 JSR and defines a programming model and architecture for implementing Web services in Java, building on J2EE, EJBs, servlets, and JSPs. The expert group work included some deep questions about Web services, such as whether Web services have homes, and what is the life-cycle model for a Web services implementation.

JSR 261

Java API for XML-based Web Services Addressing (JAX-WSA) defines an API and framework that supports addressing of Web services and messages across disparate hardware and operating systems in a transport-neutral manner.

The work in this JSR will aim to support the work done under the WS-Addressing Working Group at the W3C, <http://www.w3.org/2002/ws/addr/>. The starting point for this W3C working group was a member submission by Sun, BEA, IBM, Microsoft, and SAP to the W3C in the summer of 2004.



To explore the clickable Java Technology Concept Map, go to <http://java.sun.com/developer/onlineTraining/new2java/javamap/intro.html> or download a PDF of the Java Concept Map: http://java.sun.com/developer/onlineTraining/new2java/javamap/java_Technology_Concept_Map.pdf. The map connects the whole Java ecosystem.

“The Java platform is the most secure, reliable, and supported architecture for Web services and a service-oriented architecture”

“JAX-WSA will align closely with JSR 224 (Java API for XML-based RPC 2.0) to define the API and framework,” says Arun Gupta, who shares spec lead duties with Mike Grogan. “The call for Expert Group nominations is currently open. If you’re interested, visit the JSR 261 page for details.”

JSR 182

JPay - Payment API for the Java Platform is an interface to a J2EE Connector Architecture (JCA)–compliant Payment Resource Adapter. “It will provide a J2EE applica-

tion, such as an MP3 download server or any other Web shop or similar application, access to payment functionality via credit card, telephony account, or any other payment instrument,” says Karsten Luetzge, spec lead, Sun Microsystems. “We finished the Early Draft Review with good feedback in December, and we will update the spec accordingly. We’re now working on the Reference Implementation and Technical Compatibility Kit to finalize the JSR.”

Summary

That’s it for now. There are many more JSRs being developed by JCP members, in addition to the ones I have highlighted, but I hope you’ve enjoyed reading about some of these JSRs and that they accelerate and power your Web services development efforts. I hope to meet many of you at the Web Services Edge 2005 East Conference. If you can’t make it to the conference, you can follow the progress of these JSRs on <http://jcp.org>. As always, I am very interested in your feedback. Please e-mail me with your suggestions, opinions, and ideas. ☺

The Two Executive Committees

Who oversees the JCP program?

Two Executive Committees (ECs) are chartered to guide the development and evolution of Java technologies through the JCP program. EC members are elected by the JCP membership, and elections are held annually.

Why are there are two ECs?

There’s one to oversee the desktop/server space and one for the consumer/embedded space. The Standard/Enterprise Edition (SE/EE) EC is responsible for the J2SE and J2EE specifications, and the Micro Edition (ME) EC covers the J2ME specifications.

How many members does each EC have?

Each EC has 16 members: 10 Sun nominees ratified by JCP members and five elected members from self-nominated companies and organizations. Elections took place on November 30, 2004, adding Nortel, Google, JBoss, and Intel to the SE/EE EC, and NTT DoCoMo, Samsung, and Orange France to the ME EC. The length of an EC term is three years.

Java Web Services Developer Pack 1.5 Now Available!

The Java Web Services Developer Pack (Java WSDP) is a free integrated toolkit to help you build, test, and deploy XML applications, Web services, and Web applications.

The new Java WSDP 1.5 release contains the latest Web services technologies and standards implementations already developed and finalized through the JCP program, including XML Web Services Security, a preview of the Sun Java Streaming XML Parser based on JSR 173, and updates to existing Web services technologies previously released in the Java WSDP. It also includes guidelines for developing client-side Web services.

To download it, visit <http://java.sun.com/> and go to the Java Web Services Developer Pack home page: <http://java.sun.com/webservices/jwsdp/index.jsp>.

Got Community?

Check out the Java Web Services and XML Community or the JSR Community at <http://community.java.net/java-ws-xml/> or <http://community.java.net/jsr/>.

At these sites you’ll find people to converse with and projects such as blogs, forums, and mailing lists. The world of java.net is 112,000 members strong. Its 18 developer communities are populated with engineers, enthusiasts, evangelists, educators, industry associations, and vendors, and each visit will offer something new.

We’ve got problems with your name on them.

At Google, we process the world’s information and make it accessible to the world’s population. As you might imagine, this task poses considerable challenges. Maybe you can help.

We’re looking for experienced software engineers with superb design and implementation skills and expertise in the following areas:

- high-performance distributed systems
- operating systems
- data mining
- information retrieval
- machine learning
- and/or related areas

If you have a proven track record based on cutting-edge research and/or large-scale systems development in these areas, we have brain-bursting projects with your name on them in Mountain View, Santa Monica, New York, Bangalore, Hyderabad, Zurich and Tokyo.

Ready for the challenge of a lifetime?
Visit us at <http://www.google.com/jdj> for information. EOE





Joe Winchester
Desktop Java Editor



Software Testing Shouldn't Be Rocket Science

Earthdate: October 15, 1997, and the Cassini spacecraft is launched. Mission: to boldly go and explore the planet Saturn.

Saturn is about 10 times farther away from the Sun than the Earth, and to get there required two orbits of the inner solar system, receiving gravitational kicks from Venus and Earth before doing a flyby of Jupiter to get a final assist toward Saturn.

Piggy-backed to Cassini was the Huygens probe that would be dropped onto Saturn's moon, Titan. Unlike most other moons in the solar system that are barren, cratered rocky places, Titan has an atmosphere covering it. The purpose of the probe was to parachute through this, capturing data as it descended onto the planet's surface. The data would be transmitted from the probe up to the Cassini craft, which would act as a relay and transmit back to earth where the experiments' results would be analyzed.

On January 14, 2005, Huygens successfully landed on Titan's surface and provided some fantastic pictures of the moon (www.esa.int/SPECIALS/Cassini-Huygens/index.html). Despite this, there were two major problems on the mission.

The first is that one of the radio channels that the Huygens craft was going to use to transmit data to Cassini failed. The remaining channel was used successfully, although due to this problem only half of Huygens's pictures have come back and some experiments have had all their data lost. The reason for the problem is described as a "software commanding error." The reality is the receiver on Cassini was never programmed to switch on.

The second problem is related to the premise that Huygens transmits its collected data to the Cassini orbiter, which then relays it back to earth. Three years after the launch one of the space agency's employees became uneasy about the fact that this feature hadn't been tested enough in realistic

conditions. The story is described in detail at www.spectrum.ieee.org/WEBONLY/publicfeature/oct04/1004titan.html and provides a sobering lesson in the importance of testing. This employee worked hard to convince colleagues and superiors of the importance of testing the link in real conditions, so a simulation was done by sending data from earth to Cassini mid-mission while it hurtled toward Saturn. This mimicked the separation conditions that would be encountered between the craft and Titan and the raw data sent was echoed back to earth by Cassini and analyzed. It showed a fundamental flaw.

Because of the difference in the relative speeds at which Cassini was traveling in space relative to Huygens, there was a Doppler shift. A Doppler shift is when waves are effectively compressed if the receiver and source are moving toward each other and expanded if moving apart. As the wavelength decreases, the frequency increases, meaning that Cassini would have to adjust its listening frequencies to account for its velocity relative to the Huygens transmitter. In addition, the decoding would be affected. Digital data is split into ones and zeros and compared against a base signal to decipher; however, the Doppler shift would stretch and compress the lengths of the payload actual bits in the wave, meaning the digital signal couldn't be analyzed correctly. A fix was required to rescue the \$3.26 billion project.

Despite the fact that Cassini's hardware allows its receiver to receive over a range of shifted frequencies, the firmware program was unable to be modified after launch, even though a small fix would have sufficed. The solution they used was to alter the trajectory of Cassini's orbits of Titan so that the craft's approach allowed the radio transmissions to travel perpendicular to its direction of motion, thereby reducing the Doppler shift.

The cost of the two Cassini bugs is huge. Coming down to earth it provokes questions about testing in general. A trait I've encountered at times in my career is for a program to be released knowing it is flawed because the programmer hopes to release the working version in a subsequent fixpack, hopefully before the user has encountered the errant feature. Upgrading releases is easy for developers, but for a user who has to migrate data and schedule business downtime it's frustrating and must contribute to the perception that the latest release is not a set of fully baked features but a rollup of the previous version's fixes bundled with new features that perpetually introduce their own set of bugs.

Good testing is about attitude, where a developer takes pride not just in the elegance or volume of his or her code, but in whether it meets the user's requirements and performs reliably in its first incarnation. I once heard a developer say that releasing buggy code was part of agile programming to allow you to have more cycles of code/release//fix/code. Apart from not grasping the methodology maturely, it showed a basic lack of pride in their work that they were trying to justify. The same excuses can also lead to bloatware, where code is thrown upon code without any tight design or following the basic principles of software engineering. Is the problem that some developers are incapable of taking pride in the complete quality of their work, or that education and teaching, or marketing pressures in a commercial environment still mean that buggy software is released. Next time you release code without proper testing, keep in mind the Cassini programmers who had to physically alter their craft's passage to find a solution. While we might have the next fixpack or release available to us, in space no one can hear your excuses. ☘

Joe Winchester is a software developer working on WebSphere development tools for IBM in Hursley, UK.

joewinchester@sys-con.com

LINUXWORLD[®]

CONFERENCE & EXPO

CONFERENCE: February 14 – 17, 2005 EXPO: February 15 – 17, 2005

HYNES CONVENTION CENTER • BOSTON, MA



Linux. Reliability. Efficiency. Cost Savings. If you're thinking about Linux or Open Source, the time is now. And LinuxWorld Conference & Expo is where you need to be. Expect the following cutting edge insights:

- Forward-thinking Keynote presentations by industry thought leaders from Novell, HP, Computer Associates, and MySQL.
- 10 hours of **FREE** educational programming on everything from the State of Linux and Open Source technology to Linux and desktop users.
- The world's leading hardware and software vendors as well as the hottest new start-ups giving you the opportunity to evaluate and test drive the latest technology.

And that's only the beginning. LinuxWorld Conference & Expo – where professionals come to do business better.

Register today at linuxworldexpo.com for your **FREE** Exhibit Hall pass.

LINUXWORLDEXPO.COM

Where
OPEN MINDS
Meet



Register Online With Priority Code: D2801

PLATINUM SPONSORS



Hosting ActiveX/.NET Controls in a Java Application...



...the direct way

by Stanley Wang

This article describes how to host an ActiveX/.NET control in a Java application that is targeted for the Microsoft Windows Platform. I'll assume you know the fundamentals of Java, C++, JNI (Java Native Interface), Win32, COM (Component Object Model), and ATL (Active Template Library).

Because a .NET control can be exposed as a COM component by using the COM interoperability services provided by the .NET Framework, I'll start by hosting the ActiveX control. I'll be using the control, the ActiveX control, and the COM component interchangeably.

The ActiveX Control and Its Container Review

In short, ActiveX controls are COM controls that implement a set of standard COM interfaces. An ActiveX control is hosted by a container, which must provide a window to act as the parent for the child control and implement a set of COM interfaces for communication between the control and the container. Figure 1 shows the major interfaces implemented by the container and the control. (Please refer to the Microsoft Platform SDK for more information about ActiveX control.)

Creating the ActiveX Control and Its Container in C++

As shown in Figure 1, an ActiveX control container needs to implement a lot of interfaces. Creating such a container could be very complex. Luckily ATL provides a CAXWindowT class template that simplifies this process. (For more information about ATL, please refer to the book, *ATL Internals*.) Listing 1 shows how easy it is to create a control and its container by using ATL (all the error-checking code in this article is omitted for clarity).

CJavaCOMBridge is a class derived from an ATL CWindowImpl class template, which provides its derived classes with Windowing functionality. CJavaCOMBridge is defined in Listing 2.

All the dirty work of creating the container and control is done by axwnd_, whose type is CAXWindow, which is one of the greatest classes from ATL.

ActiveX Control Java Wrapper, Hooking an ActiveX Control to an AWT Native Interface

Adding an ActiveX control to any AWT/Swing container should be as easy as adding a button or other standard widget. To achieve this convenience, a Java wrapper, say the JAxControl, is needed. It will work as the proxy of the ActiveX control on the Java side. Naturally the JAxControl should extend from an AWT/Swing component and in one of its native methods Listing 1 will be called.

However, in Listing 1 there is a HWND typed parameter parentWnd, which is the parent window of the control container. Where does this native window handle come from?

Remember that for any AWT/Swing heavyweight component, there is a native peer associated with it. So the aforementioned native window handle could be an AWT/Swing heavyweight component's native window and it is. How to obtain it? Until Java 1.5 the only documented way to get a native window handle was through the java.awt.Canvas class (please refer to `javt.h` file, which is under the `%jdkhome%\include` directory). Listing 3 shows the relevant code.

This means that the JAxControl has to be a derived class of `java.awt.Canvas` and Listing 1 will be modified as shown in Listing 4.

Though the native handle can be obtained anytime after the Canvas is realized, a reasonable place is the `addNotify` method in which the native peer will be created. Then the JAxControl could be something like Listing 5.

Unfortunately Listing 5 won't work. The `addNotify` method is executed in the Main-Method thread while the Canvas native window is created in the Toolkit thread in which the



Stanley Wang is a software development engineer at Amazon.com. Before he joined Amazon, he was a lead software developer at Vcom3D, Inc. He has an MS degree in computer science from the University of Florida. He is interested in system programming, database systems, and networking.

stanleyycwang@yahoo.com

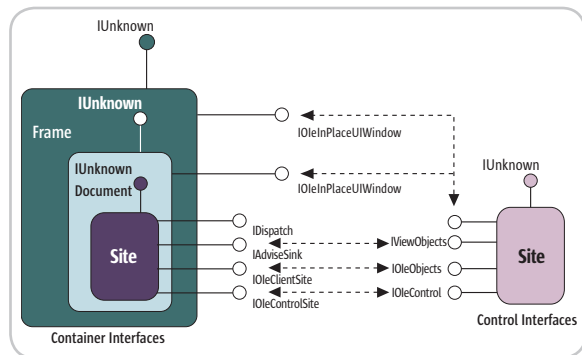


Figure 1 Interfaces implemented by the container and the control

native Windows message loop is served. (See below for more discussion about Sun Hotspot JVM threads.) This means the control created this way is useless because it can't process Windows messages. (For more information about Windows programming, please see *Win32 Programming*.) To fix this problem, the control could be created in its own native thread. Listing 6 shows one way to do it. (Listings 6–13 can be downloaded from www.sys-con.com/java/sourcec.cfm.)

Please notice that the native thread should be attached to the existing JVM to enable it to work with the JVM's synchronization scheme.

This approach solves the messaging problem, however, it has several drawbacks. First, it increases the interthread communication. Second, it violates the basic Win32 GUI programming rule because the control is created in a different thread from its parent. This could easily cause deadlock.

A better approach is subclassing an AWT Canvas native window class. A native window class is just a WNDCLASSEX structure containing all the information to describe a window. The most important member of WNDCLASSEX is the window procedure function (often called WndProc) pointer. A window procedure function is used to handle Windows messages. Remember that most of the Win32 APIs are written in C, which is not an object-oriented language. There is no Win32 API to let you derive a window class from an existing one. However, it does use object modeling to some extent. It provides some sort of polymorphic behavior for a window class by replacing its window procedure function pointer with the subclass's function pointer. If the subclass function doesn't handle a Windows message, it will pass the message to the super-class function. This is the so-called subclassing.

To subclass an AWT Canvas window class (it's named as SunAwtCanvas, however, the class name is not important here), a native subclassNativePeer method is added to the JAxControl class. Listing 7 shows the magic.

Note how a user-defined Windows message is used to create the control in the Toolkit thread. (For more information about how to define a user-defined Windows message, please see "Message Management.")

Subclassing the AWT Canvas native window class opens the back door to a mixed usage of the native window and a Java AWT/Swing component. It is efficient and powerful. The only downside is that it's relying on the AWT implementation detail.

Garbage Collection and Resource Management

JAxControl contains native resources by wrapping the ActiveX control. Though any JAxControl instance will be garbage collected eventually, it doesn't guarantee that the native resources will be freed. To avoid a memory/resource leak, special care has to be taken. A common mistake is to abuse the finalizer to reclaim resources. Unfortunately this is the wrong approach in general.

The solution is to provide an explicit termination method. Here, I'm applying the Dispose pattern that the .NET Framework has promoted, i.e., all the classes that intend to manage resources specially should implement the IDisposable interface, which has a dispose method. The dispose method will clean all the native resources, such as close a database connection and terminate child threads. Actually this pattern is applied internally by Sun; java.awt.Window, java.awt.Graphics, and some other classes all have a dispose method. Unfortunately this pattern is not abstracted to a higher level.




Figure 2 ActiveX control and its Java Wrapper

In addition to applying the Dispose pattern, it's necessary to override the removeNotify method. removeNotify is called by the toolkit internally to destroy the native peer. The ActiveX control and other native resources need to be reclaimed before the native peer is destroyed.

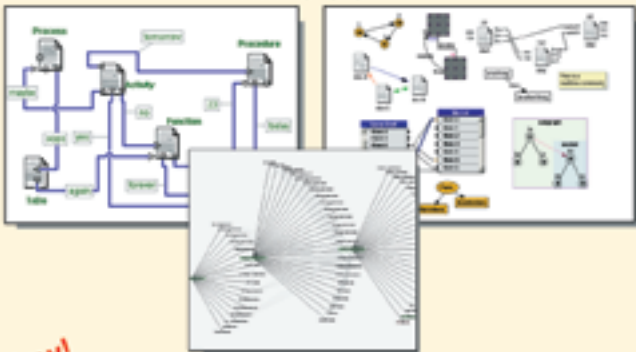
To summarize, the Java ActiveX control wrapper should look similar to Listing 8. The ComponentListener is used to adjust the control's size accordingly.

Invoking ActiveX Control Methods

With one-to-one mapping, we can create Java interfaces for any interfaces that the ActiveX control supports and that we're interested in. The control wrapper needs to implement similar interfaces as its counterpart does. Figure 2 shows the relationship between the control wrapper and the control.

One advantage to hosting the ActiveX control directly is that the control wrapper doesn't need to expose all the interfaces or all the interface methods. For example, the control in Figure 2 supports both IFoo and IBar interfaces. IFoo has two methods: foo1 and foo2. However, the control wrapper does

Build Incredible Interactive Diagrams with JGo™




New!
JGo for SWT/Eclipse
JGo Instruments for meters, dials, gauges

Create custom interactive diagrams, network editors, workflows, flowcharts, and design tools. For web servers or local applications. Designed to be easy to use and very extensible.

- **Fully functional evaluation kit**
- **No runtime fees**
- **Full source code**
- **Excellent support**

Free evaluation at:
www.nwoods.com/go
800-434-9820 or 603-886-9173



not support IBar and only supports IFoo partially. This is very useful, especially since the ActiveX control has fat interfaces.

ActiveX controls always provide access to their methods via a dispinterface and IDispatch. This is the only way for scripting languages such as JavaScript to access the control's functionality. The procedure to invoke control methods through IDispatch for C++ client is boilerplate. Listing 9 shows this procedure.

As shown in Listing 9, working with IDispatch is not efficient because each time a method ID (termed DISPID) has to be obtained first.

Fortunately most of the popular ActiveX controls support a dual interface, which means smart clients like C/C++ can access the control's custom interface directly. This is another advantage to hosting the control directly because most available commercial products only work with IDispatch. The following one line of code shows how neat, clear, and natural it is to work with a dual interface:

```
itf->foo(param);
```

Data Type Mapping

An important restriction to working with dispinterfaces and dual interfaces is that every single parameter of every single method must be variant compliant. A variant is a discriminated union and Java had no such concept until Java 1.5. Table 1 shows the core variant-compliant IDL types that map to Java.

For the primitive types there could be a direct mapping. However, some types, which have question mark entries in Table 1, have no such luck.

There are several ways to solve this problem, such as creating Java types that contain an opaque pointer to a native Variant type. However, this discussion is a topic for another day.

Cross-Apartment Access

In Listing 9 and the previous snippet, the interface pointer (disp and itf) was used directly. However, in most cases this is wrong. Generally, a windowed ActiveX control is living in an STA (single thread apartment); its interface pointer will have no meaning in another apartment. (Apartment is an execution context that defines a logical grouping of objects that share a

common set of concurrency and reentrancy constraints. For more information about the COM Apartment, please refer to *Essential COM*.) Figure 3 shows the cross-apartment access mechanics.

The question is how to get the control interface pointer in the JNI thread, i.e., how to implement the CJavaCOMBridge::QueryControl method shown in Figure 3. A common mistake is to cache an interface pointer in a static or global variable and use it in any context.

One technique for passing the interface pointer across the apartment boundary is to use interface marshaling, which will transform the interface pointer into a transmissible byte stream. An interface pointer can be marshaled by using low-level COM APIs such as the CoMarshalInterface/CoUnmarshalInterface, or by using the GIT (Global Interface Table). GIT is a process-wise singleton that contains the marshaled interface pointer that can be efficiently unmarshaled multiple times within the same process.

By using GIT, the CJavaCOMBridge class can be modified as shown in Listing 10.

There is still one thing left though. Prior to invoking any control method in a native thread, the native thread has to have an associated COM apartment. To enter a COM apartment, one of the CoInitialize/CoInitializeEx/OleInitialize family functions in the native thread can be called, depending on the apartment model and drag/drop support. Call CoUninitialize/OleUninitialize to exit the apartment before the native thread terminates. Failing to do this may delay the resources reclamation.

Sun Hotspot Java Threads Under the Hood

Because JNI functions can be called from an arbitrary thread, it's difficult to enter and exit a COM apartment from a JNI function directly. A practical approach is to do this in Java code. Before talking about this approach, let's review Sun Hotspot Java threads.

Sun Hotspot Java threads use native threads. When a Java GUI application starts, in addition to the Main-Method thread that executes the main method, there are several other threads running behind the scene (see Table 2).

In Table 2, the Reference Handler Thread and Finalizer Thread are used for garbage collection. They don't need to live in a COM apartment. The same applies to the other AWT/Daemon helper threads. However, for the rest of the threads, they need to have COM apartments associated with them if they want to use COM.

A common strategy is to wrap COM initialize and termination functions into a Java utility class and each thread finds a way to call these functions. Listing 11 shows such a COM wrapper.

Please note that there is a ThreadLocal singleton, thrdlocal, in the JCOM class. It's used to guarantee that calls to initialize/uninitialize in the same thread will be idempotent.

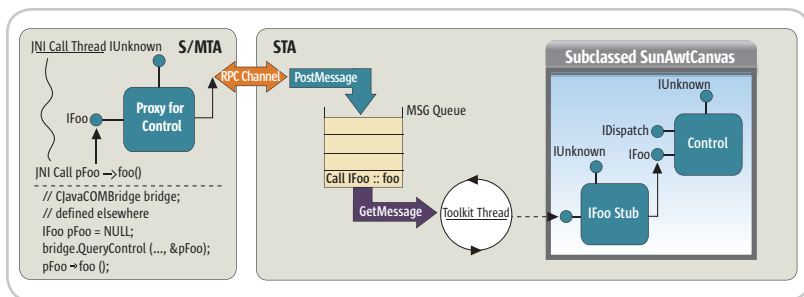


Figure 3 COM cross-apartment access mechanics

Variant IDL Type	Java Mapping	Variant IDL Type	Java Mapping	Variant IDL Type	Java Mapping
VARIANT_BOOL	boolean	float	float	IDispatch*	?
short	short	double	double	IUnknown*	?
int	int	BSTR	java.lang.String	VARIANT	?
long	long	DATE	?	SAFEARRAY	?

Table 1 Data type mapping between core variant compliant IDL types and Java types

It's easy to add initialize/uninitialize calls to the user threads and Main-Method thread, but how to inject CoInitialize/CoUninitialize calls to the Toolkit thread? Again user-defined Windows messages can be used here. Because CJavaCOMBridge subclasses the Canvas native window class, we can simply send two messages to the bridge and the bridge message handler can enter and exit a COM apartment.

How about the EDT (Event Dispatching Thread)? Remember there is a single thread rule for Swing applications: Swing components can only be accessed by one thread and generally this thread is the EDT. All the event handler functions and drawing codes will be executed in EDT. What are the good places to inject the initialize/uninitialize calls? It's not a good idea to call initialize/uninitialize in any event handler function because of the obvious order issue. It's not efficient to put these calls in a paint method because paint is called too often while the initialize/uninitialize only needs to be called once.

The question now is a general one: How do you talk to EDT from other threads? The answer is to use the SwingUtilities.invokeLater method. invokeLater will request some codes to be executed in the EDT, and its sibling, invokeAndWait, will wait for the code to be finished in the EDT. As a general rule, don't use invokeAndWait because it will cause deadlock very easily.

Now the code in Listing 12 can be added to the JAxControl. Bingo! One of the most complicated problems in COM finally has been solved.

Sink ActiveX Control Events

The ActiveX control uses an event-driven model similar to what Java does. It's using the Observer (Publisher-Subscriber) design pattern that observers can subscribe to a publisher (the control), which will call back to the observer when a state has changed (event fired).

To sink a control event in C++, you can create a regular COM object that will implement the callback interface. (This object is traditionally called a sink object and the callback interface is also called outbound interface.) The sink object doesn't need to be exportable from a DLL. Listing 13 shows the relative code.

To catch control events in Java, a common procedure can consist of the following:

- Define a Java event interface that is the mirror of the control callback interface. The Java interface does not need to contain all the methods in the control callback interface.
- Create a concrete class, say JCallbackHandler, that implements the event interface.
- Add native methods to the Control Java Wrapper, i.e., JAxControl, that will hook and unhook the C++ sink object to the concrete class instance and to the control.
- Register the concrete class instance as a listener to the Control Java Wrapper.
- When control fires an event, it will call the C++ sink object, which will in turn call back the Java concrete class instance method.

The UML sequence diagram in Figure 4 shows the procedure.

Error Handling

Virtually all COM methods return an error code of type HRESULT. HRESULT is a 32-bit integer that provides information about what type of error may have occurred. COM components that support at least an IError-Info interface are often called COM exception objects. COM exceptions can be caught by using COM exception APIs, such as GetErrorInfo().

There are generally three ways to handle COM errors and exceptions in native methods.

1. Return an error code or flag to indicate if the method has been called successfully.

Advertiser	URL	Phone	Page
Altova	www.altova.com	978-816-1600	4
blog-n-play.com	www.blog-n-play.com	201-802-3000	47
Common Controls	www.common-controls.com	+49 (0) 6151/13 6 31-0	35
DataDirect Technologies	www.datadirect.com/jdj	800-876-3101	9
EV1 Servers	www.ev1servers.net	800-504-SURF	27
Fitech Laboratories, Inc.	www.fitechlabs.com/grid	646-495-5076	33
Google	www.google.com/jdj	650-623-4000	39
ILOG	jviews-info-kit.ilog.com	800-for-ILOG	15
Intel Developer Forum	www.intelidf.com/reg4	888-261-0971	61
JDJ	www.sys-con.com/jdj	888-303-5282	49
Jinfonet	www.jinfonet.com/jp2.htm	301-838-5560	29
LinuxWorld Conference & Expo	www.linuxworldexpo.com	508-424-4810	41
M7	www.m7.com/d7.do	866-770-9770	19
Microsoft	msdn.microsoft.com/visual		11
Nexaweb	www.nexaweb.com		Cover IV
Northwoods Software Corp.	www.nwoods.com/go	800-434-9820	43
Parasoft Corporation	www.parasoft.com/jtest_jdj	888-305-0041	13
Promind	www.xmlmaker.com	415-359-9117	23
ReportingEngines	www.reportingengines.com/download/f1ere.jsp	888-884-8665	17
SAP	sdn.sap.com		21
Software FX	www.softwarefx.com	800-392-4278	Cover III
WebApp Cabaret	www.webappcabaret.com/jdj.jsp	866-256-7973	31
WebRenderer	www.webrenderer.com	+61 3 6226 6274	7
Web Services Edge 2005 East	www.sys-con.com/edge	201-802-3066	51-60
Zero G	www.zerog.com	415-512-7771	Cover II

General Conditions: The Publisher reserves the right to refuse any advertising not meeting the standards that are set to protect the high editorial quality of *Java Developer's Journal*. All advertising is subject to approval by the Publisher. The Publisher assumes no liability for any costs or damages incurred if for any reason the Publisher fails to publish an advertisement. In no event shall the Publisher be liable for any costs or damages in excess of the cost of the advertisement as a result of a mistake in the advertisement or for any other reason. The Advertiser is fully responsible for all financial liability and terms of the contract executed by the agents or agencies who are acting on behalf of the Advertiser. Conditions set in this document (except the rates) are subject to change by the Publisher without notice. No conditions other than those set forth in this "General Conditions Document" shall be binding upon the Publisher. Advertisers (and their agencies) are fully responsible for the content of their advertisements printed in *Java Developer's Journal*. Advertisements are to be printed at the discretion of the Publisher. This discretion includes the positioning of the advertisement, except for "preferred positions" described in the rate table. Cancellations and changes to advertisements must be made in writing before the closing date. "Publisher" in this "General Conditions Document" refers to SYS-CON Publications, Inc.

This index is provided as an additional service to our readers. The publisher does not assume any liability for errors or omissions.

2. Convert the error code into a Java exception, which will be caught in Java code.
3. Throw a Java runtime exception.

Depending on the requirements, one of these approaches or a combination of them can be applied.

Windowless ActiveX Control

So far I have only talked about the Windowed ActiveX control. However, some ActiveX controls, or more generally speaking COM components, don't have a window. It's much easier to host these controls. You can still follow the procedure I described in this article, but the Control Java Wrapper doesn't need to extend from Canvas anymore and can be hosted in the Toolkit thread.

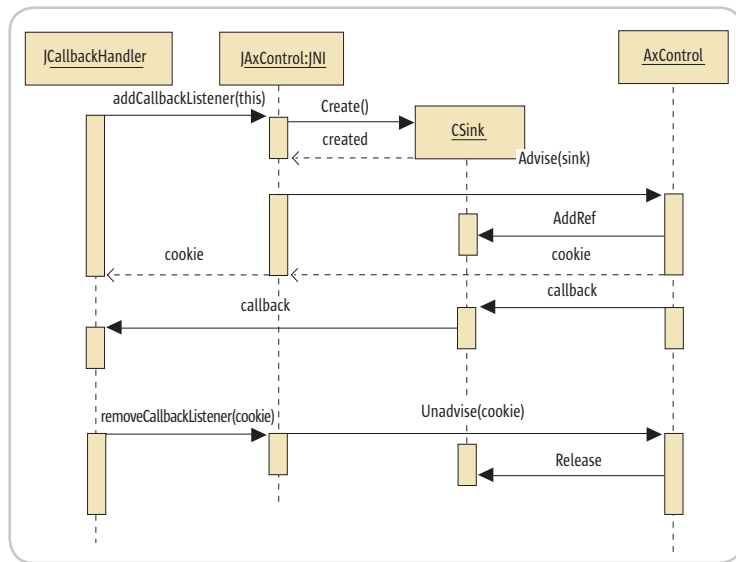


Figure 4 Sink control events

For Windowless controls, they may live in an MTA (multithreaded apartment). If this is the case, it's better for the hosting thread to enter the MTA for best performance.

.NET Control

Armed with the above knowledge, it will be a breeze to host a .NET control by using the COM interoperability services provided by the .NET Framework.

After a .NET control such as a Windows Form control has been created, the first step to making it creatable in a control container is to export .NET types. Remember that the most common way to activate a COM object is to call the CoCreateInstance API, which will look in the registry for the dll server location that provides the COM object. A .NET control can be registered in a registry by using a Visual Studio IDE or a RegAsm tool. RegAsm will search the metadata of the assembly passed to it for any public creatable .NET types. For each suitable type it will generate COM registration information.

Notice that the InprocServer32 key is specifying mscoree.dll as the server, which is different from the traditional COM object whose InprocServer32 key is specifying the server dll. The .NET type is contained in the assembly dll, which is given by the CodeBase value. If there is no CodeBase value specified, the assembly dll needs to be shared by putting it into the GAC (Global Assembly Cache). The Class value specifies the .NET type that is exported.

The next step in using the .NET control is to import metadata in the source file into the JNI project. This can be easily done by using the #import directive. The following line will import the type information:

```
#import "jtlnet.dll" raw_interfaces_only named_guid
```

After this, the .NET control can just be treated as a regular COM control.

JMediaPlayer Example

JTL, which is an open sourced, JNI C++ template library that I've written, provides a complete example of how to host the Microsoft Media Player into a Java application. The source code can be downloaded from <http://sourceforge.net/projects/jtl/>.

Summary

In this article I described how to host an ActiveX/.NET Control directly in a Java application. I also presented several tricks for solving Java's nondeterministic destruction issue, subclassing the Canvas native window class, creating COM cross-apartment access, and injecting code into the Sun Hotspot JVM threads. Though there are still some topics left, such as complex data type mapping, sinking control property notification events, and accelerator translations, the methods I showed in this article can be easily extended to support them. If you need more functionality, jtl::win::javacom is a good framework to start with.

Acknowledgment

Special thanks to Bino George at the Sun AWT Team, especially for his reviews.

Java Threads	Responsibility
Reference Handler Thread	Keeping track of the object references
Finalizer Thread	Dispatching finalize calls on non-accessed objects
AWT Event Dispatching Thread (EDT)	Monitoring system EventQueue, dispatching events, handling repaints
AWT Toolkit Thread	Implementing the native WndProc window procedure function, converting Windows events to Java events, posing the Java events to the EventQueue
Main-Method Thread	Executing application's main method
User Thread	User launched thread (a.k.a Runnable or Thread)
Other AWT/Daemon Threads	...

Table 2 Sun Hotspot Java threads




Figure 5 A .NET control registration information

References

- Microsoft Platform SDK: www.microsoft.com
- Rector, B., and Sells, C. (1999). *ATL Internals*. Addison-Wesley Professional.
- Rector, B., and Newcomer, J. (1997). *Win32 Programming*. Addison-Wesley.
- Newcomer, J. (2000). "Message Management": www.codeproject.com/dialog/messagemgmt.asp
- Bloch, J. (2001). *Effective Java Programming Language Guide*. "Item 6: Avoid Finalizers." Addison-Wesley Professional.
- Box, D. (1998). *Essential COM*. Addison-Wesley Professional.
- Gamma, E., et al (1995). *Design Patterns*. Addison-Wesley Professional.
- Troelsen, A. (2002). *COM and .Net Interoperability*. Apress.
- Wang, S. (2004). "Embedding the Java Virtual Machine, Once and for All." *JDJ*, Vol. 9, issue 8.

Listing 1

```
// cpp file
CJavaCOMBridge bridge
bool createControl(HWND parentWnd) {
    ...
    bridge.Create(parentWnd, CWindow::rcDefault);
    ...
}
```

Listing 2

```
// cpp file
class CJavaCOMBridge : public CWindowImpl<...> {
private:
    CAXWindow axwnd_;
public:
    ... // Message Map
    LRESULT OnCreate(...){
        ...
    // create container
    axwnd_.Create(m_hWnd, CWindow::rcDefault);
    // create ActiveX control
    axwnd_.CreateControl(progid);
    ...
    }
    ... // other stuff
}
```

Listing 3

```
HWND GetNativeWindow(JNIEnv* env, jobject canvas, jint awtversion)
{
    // Get the AWT
    JAWT awt;
    awt.version = awtversion;
    JAWT_GetAWT(env, &awt);
    // Get the drawing surface
    JAWT_DrawingSurface* ds = awt.GetDrawingSurface(env, canvas);
    // Lock the drawing surface
    jint lock = ds->Lock(ds);
    // Get the drawing surface info
    JAWT_DrawingSurfaceInfo* dsi = ds->GetDrawingSurfaceInfo(ds);
    // Get the platform specific drawing info
    JAWT_Win32DrawingSurfaceInfo* dsi_win =
        (JAWT_Win32DrawingSurfaceInfo*)dsi->platformInfo;
    return dsi_win->hwnd;
}
```

Listing 4

```
// jni cpp file
#include <jniutil.hpp> // defined GetNativeWindow
JNIEXPORT jboolean JNICALL Java_JAxControl_createControl(JNIEnv* env, jobject canvas) {
    ...
    HWND parentWnd = GetNativeWindow(env, canvas, awtversion);
    bridge.Create(parentWnd, CWindow::rcDefault);
    ...
}
```

Listing 5

```
// java file
public class JAxControl extends Canvas {
    public JAxControl() {
        ...
    }
    public void addNotify() {
        super.addNotify();
        createControl();
    }
    ... // other stuff
}
```

THREE REASONS TO

blog-n-play.com

1 Get instantly published to 2 million+ readers per month!

blog-n-play™ is the only FREE custom blog address you can own that comes with instant access to the entire i-technology community. Have your blog read alongside the world's leading authorities, makers and shakers of the industry, including well-known and highly respected i-technology writers and editors.

2 Own a most prestigious blog address!

blog-n-play™ gives you the most prestigious blog address. There is no other blog community in the world that offers such a targeted address, and comes with an instant targeted readership.

3 Best blog engine in the world...

blog-n-play™ is powered by *Blog-City*™, the most feature rich and bleeding-edge blog engine in the world, designed by Alan Williamson, the legendary editor of *JDJ*. Alan kept the i-technology community bloggers' demanding needs in mind and integrated your blog page to your favorite magazine's Web site.



www.TAMI.linuxworld.com

"Many blogs to choose from"


PICK YOUR MOST PRESTIGIOUS ADDRESS

IT Solutions Guide	MX Dev. Journal
Storage+Security Journal	ColdFusion Dev. Journal
JDJ: Java	XML-Journal
Web Services Journal	Wireless Business &Tech.
.NET Dev. Journal	WebSphere Journal
LinuxWorld Magazine	WLDJ: WebLogic
LinuxBusinessWeek	PowerBuilder Dev. Journal
Eclipse Dev. Journal	

3 MINUTE SETUP

Sign up for your FREE blog Today!





Oracle Application Server 10g

Oracle has impressed most of us with its database products and, in my opinion, they offer an impressive J2EE infrastructure as well. To select the right technology for migrating the popular Alamo Rent A Car Web site from its generation-old Microsoft platform to a state-of-the-art J2EE stack, we took a close look at Oracle's latest product, Oracle Application Server 10g, in conjunction with the well-integrated Oracle JDeveloper Integrated Developer Environment. During this review, Oracle's product established itself as a formidable choice for a standards-compatible J2EE development and runtime platform. It continues to prove its enterprise-grade reliability by running the mission-critical Alamo.com Web site almost flawlessly since August. What is equally important to us is that it's a choice for the future, with industry-leading support for service-oriented architecture (SOA) and the Grid. With its Business Process Execution Language (BPEL) support for orchestrating services, Oracle Application Development Framework (ADF), integration capabilities for leveraging existing heterogeneous assets (for example, Tuxedo services written in C), a certificate authority, and even Web application service level monitoring, this product delivers a lot of bang for the buck.

Standard J2EE Support

Before talking about the bells and whistles, let's first stick to the basics. Compliance with J2EE standards was a must – we could not compromise. We extensively tested and later successfully deployed on the Alamo.com Web site all of the following J2EE 1.3 APIs:

- JavaServer Pages 1.2
- Java Servlet 2.3
- Enterprise JavaBeans 2.0

- JDBC 2.0 Extension
- 2EE Connector Architecture 1.0
- Java API for XML Processing 1.1
- Java Naming and Directory Interface (JNDI)

J2EE continues to be a key component of Oracle's corporate strategy.

Early on, we decided to leverage standard design patterns for this large-scale Web site. J2EE APIs were to be employed across all layers of the application architecture via the use of:

- Model-View-Controller (MVC) architecture
- Business delegates
- Session façades
- Service locators
- Data Access Objects

How could Oracle JDeveloper, with its ADF, aid our approach? How would it fit our standards-based development established on MVC?

We have found that Oracle's ADF helps manage the ever-increasing complexity of J2EE applications. The framework is based on an MVC architecture and allows the use of any view, model, or controller technology. The declarative data binding that Oracle developed with ADF is part of JSR (JSR-227). The framework's visual and declarative capabilities, coupled with guided-coding techniques, can increase developer productivity. We developed an application with a JSP view layer, Struts Controller, and EJB model layer using the framework. The tool is a good fit if you follow standard design patterns (we used Struts as the ADF controller), and can significantly reduce the amount of infrastructure coding.

Oracle JDeveloper is a comprehensive, integrated IDE. It also provides built-in features for standard tools such as JUnit and CVS, which are important for our development

Oracle Corporation

500 Oracle Parkway
Redwood Shores, CA 94065
www.oracle.com
Phone: 650 506-7000
www.oracle.com/technology/products/ias/index.html

and testing process. We experienced some issues in accommodating team development; however, Oracle is addressing these issues. We were happy to hear that the next release of Oracle JDeveloper implements several enhancements for team development.

Support for Mission-Critical J2EE Deployments

Our second checkbox for selecting a serious J2EE platform was support for high availability. The Alamo Web site had to be designed for high availability, which means minimizing unplanned downtime. As you would expect from an enterprise-grade player such as Oracle, Oracle Application Server 10g did not disappoint. We successfully tested and later deployed the following built-in capabilities:

- Automatic connection rerouting
- Load balancing in the instance of a failure
- Clustered session state failover across a cluster
- Fast-start fault recovery architecture where failures or instance deaths are automatically discovered and instances are automatically restarted, reconfigured, and reclustered

Furthermore, with Oracle Enterprise Manager's all-Web-based console (see Figure 1), you can manage such a high-availability configuration and monitor system resources such as CPU usage, memory, and

Mark Euler works for Perot Systems and currently serves as development manager for Vanguard Car Rental e-commerce applications, and systems architect for the Odyssey application, a large multitier OLTP that runs much of Vanguard's business. He recently led a team of 15 developers in designing Vanguard's future e-commerce architecture and redesigning the Alamo Web site, which included the development of several applications to support the Web site functionality.

mark.euler@ps.net

even URL performance (Application Service Level Monitoring).

It is worth noting that the August 2004 relaunch of the Alamo Web site built on Oracle Application Server 10g went off without a hitch. Since then, we have not experienced a single unplanned outage due to a Web component. All technology aside, that's what really counts.

A Complete Application Server Suite Package

Now we move beyond what you'd expect from a pure J2EE application server and highlight some key add-ons Oracle Application Server 10g delivers as part of its complete package: integration and identity management.


To test the integration capabilities of Oracle's application server (in the form of the Oracle Application Server Integration component), we built an application that leveraged existing business logic in C and Tuxedo. Using Oracle Integration and the adapters, our application exposes Tuxedo services within EJB components of several J2EE ap-

plications. This is not your everyday, standard J2EE requirement; nevertheless, it's not uncommon for a Web site operating in a heterogeneous setting. Oracle Application Server Integration and the accompanying adapters delivered on our requirement to leverage existing code. Because it behaves as a Tuxedo client, Oracle Integration does the work of transforming C language data structures into XML structures (defined at development time) that are subsequently transformed into domain objects in the EJB layer and used within the J2EE application. Without this capability, leveraging Tuxedo services within our J2EE applications would have been quite time-consuming. My only criticism addresses the efficiency of the process modeling-driven approach of the previous process integration tool. We found this to be a cumbersome way to manage the data structures and mappings required for process integration. Oracle BPEL Process Manager promises significant improvement over its predecessor.

Security is a key concern for most enterprise IT shops. A typical enterprise J2EE environment requires a robust security and identity management infrastructure. Oracle goes beyond the standard J2EE requirements, such as incorporating a JAAS provider. Oracle's Single Sign-On module provides access control to Oracle Application Server Portal (integrated into the application server) and applications, and is supported by Oracle Internet Directory (OID), which stores user identities and passwords. OID comes with an out-of-the-box Web-based GUI to set up users, which is part of the Delegated Administrative Service (DAS). We tested these capabilities and successfully leveraged the functionality to consolidate user access to internal applications.

Note that there are several further platform components we were unable to test as part of this review. These include TopLink (a persistence layer providing OR mapping), Oracle Application Server Portal, and Oracle Application Server Business Intelligence.

The World's Leading Java Resource Is Just a >Click< Away!



JDJ is the world's premier independent, vendor-neutral print resource for the ever-expanding international community of Internet technology professionals who use Java.

Only **\$69⁹⁹** ONE YEAR 12 ISSUES

Subscription Price Includes **FREE** JDJ Digital Edition!

www.SYS-CON.com/JDJ
or **1-888-303-5282**

OFFER SUBJECT TO CHANGE WITHOUT NOTICE

What Sets Oracle Application Server 10g Apart

Let's focus some attention on functionality details, which, in my opinion, set Oracle Application Server 10g apart from other J2EE platforms such as IBM WebSphere and BEA

WebLogic (also part of our evaluation). This functionality includes extensive caching technology, which yielded unexpected performance benefits. Although we have not had the opportunity to use the native BPEL capabilities of the BPEL Process Manager, we have

been impressed with what we've seen in our evaluations.

As for caching, we tested both Oracle Application Server Web Cache and Java Object Cache (JOC) caching technologies. Used at the network edge, Web Cache combines caching and load balancing. It intercepts HTTP requests and responds with a cached page if one is available. If one is not available, it passes the request back to the least busy application server. Our expectations for performance improvements on our dynamic e-commerce Web site were not high, but Web Cache also has the ability to cache dynamically generated, next-to-static Web content. The result: on Alamo.com, Web Cache serves 80% of all requests, of which 25% are served in compressed form – a clear performance boost!

We also reviewed JOC, which is part of the OC4J J2EE container. JOC provides an API and service for managing cache objects within J2EE apps. We truly loved this feature because it increases performance, especially in a clustered environment. Oracle contributed this innovation to the J2EE community as JSR 107 (JCache API).

Let's not end this review without talking about what Oracle sees as the "crown jewel" of the Application Server 10g platform, Oracle's next-generation process integration product, the Oracle BPEL Process Manager (see Figure 2). Because it's a standards-based BPEL engine, it is a comprehensive offering, not only for its standards-based Web services orchestration capabilities, but also for the built-in WSIF bindings that enable BPEL to orchestrate nearly any legacy system as if it were a Web service – without having to explicitly wrap or publish it as one.

SOA and Web services are certainly the future. With a native BPEL implementation in Oracle BPEL Process Manager, Oracle is among the first to offer the crucial service orchestration capability needed to build service-oriented applications easily. Coupled with a comprehensive and integrated SOA platform in Oracle Application Server 10g, Oracle is a compelling choice for realizing enterprise SOA!



Figure 1 Oracle Enterprise Manager 10g: analyzing URL performance

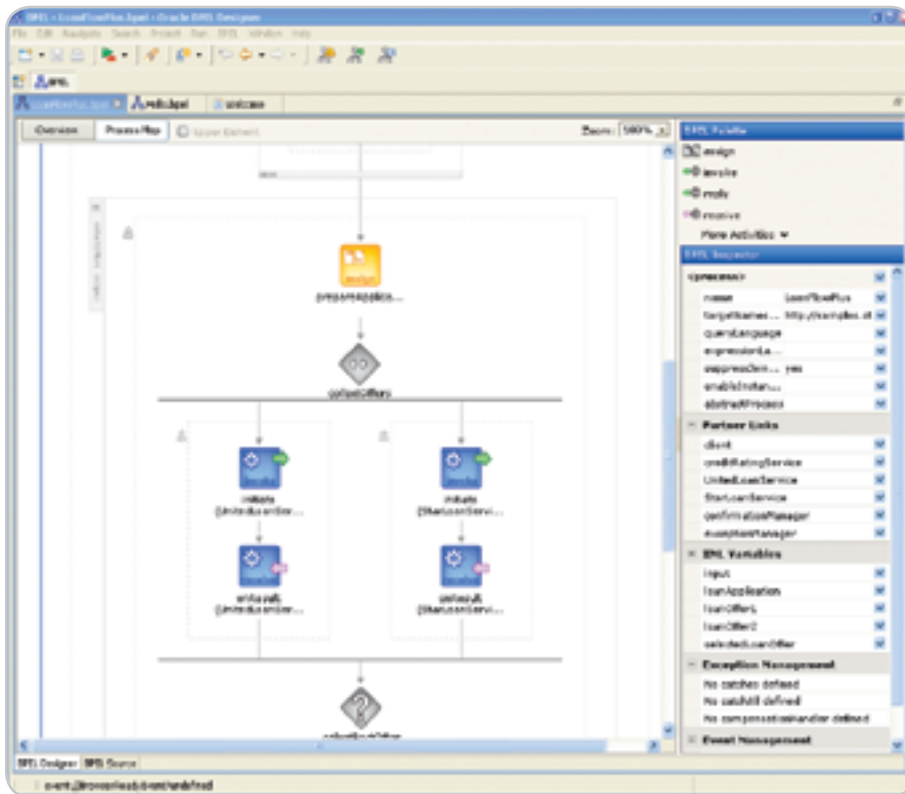


Figure 2 Oracle BPEL Process Manager – process development

See what everyone is talking about... **The Application Server Shoot-Out!**
International Web Services Conference & Expo

Keynotes



Matt Ackley
eBay



Ari Bixhorn
Microsoft



Mike Milinkovich
Eclipse.org



Register by
FEBRUARY 11, 2005
SAVE Up To
\$200

WEB SERVICES EDGE 2005 EAST

**Product Teams
 Delivering SOA,
 Web Services,
 Java, .NET,
 & XML Solutions**



New for 2005...

- Cabana Night Developer Exchange
- Application Server Shootout
facilitated by Burton Group
- Hands On VS and .NET Labs
- Software Giveaway



Free Tutorials presented by **critical sites** **Novell** **Microsoft** Java Community Process

For more information:
 visit: www.sys-con.com/edge
 call: 201-802-3066
 e-mail: jimh@sys-con.com

February 15-17, 2005
John B. Hynes Convention Center
Boston, MA

PRODUCED BY:
SYSCON MEDIA

PRESENTED BY:
SYSCON EVENTS

DIAMOND SPONSOR: **GOLD SPONSORS:** **ASSOCIATION SPONSORS:**

CORPORATE SPONSOR: **APP-SERVER SHOOTOUT FACILITATOR:**

MEDIA SPONSORS:

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

International Web Services



Web Services Edge 2005 East

February 15-17, 2005
John B. Hynes Convention Center
Boston, MA



The Largest *i*-Technology Event of the Year!

Who Should Attend

- Software Developers
- Software Engineers
- Application Developers
- Development Managers
- Technical Directors
- Analysts/Programmers
- IT Managers
- CEOs
- CTOs
- CIOs
- Technical Architects
- Team Leaders
- Software Consultants
- IT Directors
- Project Managers

Features & Attractions

- 3 Days Packed with Education and Training
- Keynotes & Panel Discussions from Industry Leaders
- 50+ Hard-hitting and Informative Seminars
- FREE Web Services Tutorial Presented by Novell
- FREE .NET Tutorial Presented by Microsoft
- FREE Web Services Security Tutorial Presented by Critical Sites
- Cabana Night Developer Exchange Presented by INETA
- Panel Discussions Presented by JCP
- Opening Night Welcome Reception
- Compelling Case Studies & Best Practices
- Hands-On Labs Featuring .NET and Visual Studio Presented by Microsoft Partners
- Featured Product Demonstrations
- Exhibit Floor featuring hundreds of products
- Real-time SYS-CON Radio Interviews

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

VISIT www.sys-con.com/edge FOR TIMES & SCHEDULES

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

ces Conference & Expo



Keynote Speakers

Tuesday, February 15, 11 a.m.

Matt Ackley

SENIOR DIRECTOR, eBay
DEVELOPERS PROGRAM



Web Services for eCommerce

eBay, The World's Online Marketplace, has more than 114 million registered users, 10,000 developers, and over 700 live, third-party applications. Four years ago, eBay began allowing third parties to build applications that tap into eBay, and today eBay hosts one of the leading Web services platforms. Through its developer program, eBay enables third parties to create cutting-edge Web services applications that benefit the buyers and sellers on eBay. At present, 40% of eBay's listings come through its API, which handles more than a billion Web services calls a month. Ackley will discuss the rewards and challenges of building and maintaining one of the world's leading Web services platforms, and share insights and practical guidelines for others.

Matt Ackley is senior director of the eBay Developers Program. He supports eBay's vision to be the leading platform for global online commerce, and is chartered with creating a thriving ecosystem between eBay, its community of users, and third-party developers and solution providers. Ackley joined eBay in 2003 as part of eBay's acquisition of FairMarket, which provided technology solutions and services to online marketplaces.

Wednesday, February 16, 11 a.m.

Ari Bixhorn

DIRECTOR, WEB SERVICES STRATEGIES,
MICROSOFT CORPORATION



Introducing Indigo: The Unified Programming Model for Building Service-Oriented Applications

Indigo is Microsoft's unified programming model for building service-oriented applications on the Windows platform. It enables developers to build secure, reliable, transacted solutions that integrate across platforms and interoperate with existing investments. Indigo combines and extends the capabilities of existing distributed application technologies, including .NET Enterprise Services, System.Messaging, Remoting, ASMX, and WSE to deliver a unified development experience spanning distance, topologies, hosting models, protocols, and security models. This keynote will provide an inside look at Indigo and show you how Indigo will radically simplify the development of distributed, service-oriented applications.

Ari Bixhorn is the director of Web Services Strategy in the Developer and Platform Division at Microsoft Corp. He is responsible for product planning and technical evangelism for Microsoft's Web services offerings, including "Indigo," the code name for a component of the next version of the Windows operating system, code-named Windows "Longhorn." Bixhorn has spent the past five years at Microsoft, driving product management efforts for the Visual Basic and Visual Studio development systems.

Thursday, February 17, 11 a.m.

Mike Milinkovich

ECLIPSE.ORG



An Open Development Platform for Web Services

Open source technology runs the Internet. Linux, Apache, PHP and Eclipse are highly successful open source communities that provide the backbone for today's Web applications. All indications point to a continued value proposition for organizations for leveraging open source when developing and deploying SOA-based applications. This keynote will examine the benefits of using open source technologies, the decision-making process used when adopting these solutions and the potential for contributing back to the open source community.

Mike Milinkovich has held key management positions at Oracle, WebGain, The Object People, and Object Technology International Inc. (which subsequently became a wholly owned subsidiary of IBM), assuming responsibility for development, product management, marketing, strategic planning, finance, and business development. Mike earned his MS degree in information and systems sciences and a bachelor of commerce degree from Carleton University in Ottawa, Canada.

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

VISIT www.sys-con.com/edge FOR TIMES & SCHEDULES

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

DAY 1
FEBRUARY 15
DAY 2
FEBRUARY 16
DAY 3
FEBRUARY 17

	Java	.NET	Web Services
7:30	Registration		
8:00	FREE Tutorial – Ashish Larivee, Novell, Using a Web Services Framework to Build SOA Applications		
9:00	(J-1) What's New In JDO 2.0	(.NET-1) Intro to SPOT	(WS-1) Ensuring Web Services Interoperability
10:00	(J-2) Using Java Messaging in Real-Time Trading Systems	(.NET-2) An Introduction to SQL Server Reporting Services	(WS-2) Web Services Standards: Going Behind the Mask
11:00	Opening Keynote – Matt Ackley, Senior Director, eBay Developer Program, eBay		
12:00	EXPO OPEN (12 P.M.–5 P.M.)		
3:00	Keynote Panel Presented by JCP – Web Services and Security Moderator: Onno Kluyt, Sr Director & Chair, JCP Program, Sun Microsystems		(WS-2B) Solving Complex Business Problems Through SOA
4:00	(J-3) The ROI of a Java-Rich Client	(.NET-3) Go With The Flow – Human Workflow Services in BizTalk 2004	(.NET-3B) Techniques with Visual Basic.NET (WS-3) The XML Data Challenge
5:00	Opening Night Reception		

	Java	.NET	Web Services
7:30	Registration		
8:00	FREE Tutorial – Thom Robbins, Microsoft – The Next Generation of Visual Studio (free with VIP preregistration)		
9:00	(J-4) Web Services End-to-End Security on J2EE: Gaps and Proposed Solutions	(.NET-4) The Microsoft Framework: An Agile Software Development Process for Building Web Service Applications	(WS-4) How To Bulletproof Your Web Services
10:00	(J-5) J2ME and Eclipse	(.NET-5) Web Services Security for Dummies with WSE2	(WS-5) The Role of Policy in Web Services Integration – It's More Than Just Security
11:00	Keynote – Ari Bixhorn, Director, Web Services Strategies, Microsoft Corporation		
12:00	EXPO OPEN (12 P.M.–4 P.M.)		
3:00	Application Server Shootout		
4:00	(J-6) The Impact of JBoss and Mono on the Application Server Market	(.NET-6) Securing Service-Oriented Architecture with Microsoft's WSE 2.0	(.NET-6B) J2EE to .NET Interoperability and App. Integration (WS-6) B2B Policy Enforcement: The Third Rail of SOA Implementation
5:00	(J-7) Migrating Enterprise Applications Between J2EE Application Servers	(.NET-7) So You THINK You Know What an Object Is...	(.NET-7B) Building and Using Advanced ASP.NET Web Controls (WS-7) Driving SOA Governance
6:00	Cabana Night – Hosted by INETA		

	Java	.NET	Web Services
7:30	Registration		
8:00	FREE Tutorial – Patrick Hynds and Duane Laflotte, Critical Sites – Security, The New Reality (free with VIP preregistration)		
9:00	(J-8) Design Patterns and Project Organizational Techniques for "Write Once, Debug Everywhere"	(.NET-8) Migrating ASP to ASP.NET	(WS-8) SOA: From Pattern to Production
10:00	(J-9) Using Grid Computing with Web Services and J2EE to Create Internet-based SOAs	(.NET-9) Smart Client Development with the Offline Application Block	(WS-9) High Performance Web Services – Tackling Scalability and Speed (WS-9B) Effective Risk Abatement and Success in a Service-Oriented World
11:00	Keynote – Mike Milinkovich, Executive Director, Eclipse Foundation		
12:00	EXPO OPEN (12 P.M.–4 P.M.)		
3:00	(J-10) Java Web Services Programming Tips & Tricks	(.NET-10) CLR Internals	(WS-10) So You Want an SOA: Best Practices for Migrating Toward Service Orientation in the Enterprise
4:00	(J-11) JCP Program: How the Java Technology Binary Software Standard is Managed and Evolves	Visit Web site for update	(WS-11) Four Abilities SOA Will Lack Without a Registry

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

WS Security	Case Study
FREE Tutorial –Aaron Williams, JCP, <i>Developing Web Services Using Java Technology</i>	
(WSS-1) Identity in SOA	(CS-1) Developing E-Commerce Applications with Web Services
(WSS-2) Securing Web Services with WS-Security	(CS-2) Developing Enterprise Class Web Services
	(CS-3) Service-Oriented Development on NetKernel – Patterns, processes and product to reduce the complexity of IT systems
(WSS-3) Anatomy of a Web Services Attack	

Novell.

FREE Web Services Security Tutorial Presented by Novell

Using a Web Services Framework to Build SOA Applications

Tuesday, Feb. 15
8:00 A.M. – 11:00 A.M.

FREE Tutorial Presented by Java Community Process

Developing Web Services Using Java Technologies

Tuesday, Feb. 15
8:00 A.M. – 11:00 A.M.



Free Tutorials with VIP Preregistration ONLY!

WS Security	Case Study
(WSS-4) Using Mobile Phones as an SSO Authentication Device in SOA Solutions	(CS-4) Orchestrating FORCEnet Engagement Packs with BPEL for Web Services
(WSS-5) Building Intelligent Enterprises with Novell's Identity-Driven Computing	(CS-5) CPI: A Globally Integrated Problem-Tracking and Resolution System Using Java Web Services
(WSS-6) XML Content Attacks	(CS-6) The Transformation of SiteRefresh into a Web Services
(WSS-7) The Interoperability Challenge of Web Services Security Standards	

Microsoft®

FREE .NET Tutorial Presented by Microsoft

The Next Generation of Visual Studio
Thom Robbins

Wednesday, Feb. 16
8:00 A.M. – 11:00 A.M.

Free Tutorials with VIP Preregistration ONLY!

WS Security	Case Study
(WSS-8) Transitioning Successfully to SOA and Web Services: Building the Infrastructure for SOA Growth	(CS-8) Using SOA and Web Services to Issue Business Licenses in the District of Columbia
Visit Web site for update	(CS-9) Developing Web Services with Eclipse



FREE .NET Tutorial Presented by Critical Sites

Security, The New Reality
Patrick Hynds

Thursday, Feb. 17
8:00 A.M. – 11:00 A.M.

Free Tutorials with VIP Preregistration ONLY!

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

International Web Services

JAVA TRACK

J-1 What's New In JDO 2.0

PATRICK LINSKEY, SOLARMETRIC

Tuesday, February 15, 2005 9:00 A.M. – 9:50 A.M.

In this presentation, SolarMetric's CTO Patrick Linskey introduces JDO, the standards-based object persistence specification, and the new changes introduced with the release of JDO 2. Developers using JDO are seeing increased application portability, reduced development cycle time, improved code quality, higher performance, and more manageable scalability. The session will highlight the problems of persisting data to a database and the basics of JDO. Patrick will focus on what's new with JDO 2, including the improvements and alternatives to the JDO Query Language, detached database operations, and the new metadata for object / relational mapping.



BIO: Patrick Linskey has been working with Java Data Objects for over 3 years and has been involved in object/relational mapping for 5+ years. As the founder and CTO of SolarMetric, Patrick drives the technical direction of the company. Patrick is a luminary on JDOcentral, a consortium committed to marketing the JDO standard. He has been one of the leaders on the JDO specification team, currently helping to drive the JDO 2.0 specification.

J-2 Using Java Messaging in Real-Time Trading Systems

YAKOV FAIN, SMART DATA PROCESSING, INC.

Tuesday, February 15, 2005 10:00 A.M. – 10:50 A.M.

Any modern financial trading system is a complex distributed application that runs on multiple platforms and consists of components and services that have to communicate with each other. This presentation is about using JMS and Message-Oriented Middleware as a backbone that provides a fast and reliable delivery mechanism between various components and tiers of such systems. This presentation is based on my real-world experience in design, development, and production support of an application that had to wire together midrange computers with J2EE and LDAP servers, non-Java stock exchange software, and mainframe legacy reporting systems.



BIO: Yakov Fain works as a Java architect for a major bank in New York City. He is a member of the editorial board of *JDJ*. Yakov has 23 years of experience working in software development.

J-3 The ROI of a Java-Rich Client

COACH WEI, NEXAWEB TECHNOLOGIES

Tuesday, February 15, 2005 4:00 P.M. – 4:50 P.M.

Rich client, thick client, thin client – what are they? What are the benefits of rich client, in particular Java Rich Client? This session presents an overview of the various client technologies, in particular various Java-based rich-client solutions, including Swing/AWT, SWT, and XML. Further, it analyzes the trade-off between these different Java-rich client technologies, and presents real-world case studies to justify the ROI of Java-Rich Client solutions.



BIO: Coach Wei is founder and CTO of Nexaweb Technologies (www.nexaweb.com), which develops the leading software platform for building and deploying Enterprise Internet Applications. Coach has his master's degree from MIT, holds several patents, is the author of many technology publications, and is an industry advocate for the proliferation of open standards.

J-4 Web Services End-to-End Security on J2EE: Gaps and Proposed Solutions

SUDHIR BHOJWANI, SUSHIL SHUKLA, & SUDHRITY MONDAL, BEA

Wednesday, February 16, 2005 9:00 A.M. – 9:50 A.M.

Even though WS-* security standards (WS-Security, WS-Trust, WS-SecureConversation, WS-Policy, etc.) are sufficiently prescriptive on specific security subjects like signing, SOAP message encryption, request/receive security tokens, they do not provide end-to-end security protocol that Web services can depend on to meet their security requirements. The most significant gap is identity propagation from a Web service into a J2EE container. Current JAX-RPC specification or JSR 109 does not cover this issue. This presentation identifies the gaps

and discusses the approaches to plug these gaps. It also discusses implementation of a solution for identity propagation from client to Web service and from Web service to J2EE container.



BIO: Sudhir Bhojwani is a senior architect working for BEA Technical Solutions Group. He has over eight years of industry experience architecting and designing systems based on component technologies like CORBA and EJBs, and most recently working on SOA principles.



BIO: Sushil Shukla is a principal architect working for BEA Technical Solutions Group. He has 19 years of industry experience and 9 years of architecting and designing application server-based large systems. He has extensive experience with tools and technologies used in developing mission-critical applications.



BIO: Sudhrity Mondal is a principal architect working for BEA Technical Solutions Group. He has 14 years of industry experience and specializes in Java technology, enterprise integration solutions, object-oriented software design, and development and software patterns. He has helped prominent BEA Customers to architect, design, and implement new systems using Tuxedo, J2EE, and Web services over the past eight years.



J-5 J2ME and Eclipse

MICHAEL VAN MEEKEREN, IBM

Wednesday, February 16, 2005 10:00 A.M. – 10:50 A.M.

Eclipse provides support for Java program development such as editing, compiling, and debugging, and is readily extensible through its plug-in mechanism. Many have been involved in the development of plug-ins that support the building and launching of embedded applications (with support for various platforms, such as J2ME/MIDP, PocketPC, and PalmOS). This talk will show how applications can be developed, compiled, analyzed, and compressed to fit on really small devices. It will include reports on practical experience, it will provide background information on developing Java applications for resource-constrained environments, and explain what Java standardization processes are under way.



BIO: Michael Van Meekeren has been a senior developer with IBM Ottawa Labs (formerly known as Object Technology International) since 1994, and has played an active role in the development of Envy/Developer, IBM Smalltalk, VisualAge for Java, and WebSphere Studio Device Developer. Michael is currently the IBM Eclipse Platform UI Team Lead at IBM Ottawa Labs.

J-6 The Impact of JBoss and Mono on the Application Server Market

PIERRE FRICKE, D. H. BROWN ASSOCIATES

Wednesday, February 16, 2005 4:00 P.M. – 4:50 P.M.

Linux and Apache drove a dramatic change in the server operating system and Web server marketplaces. These areas, dominated by Microsoft and Sun in the late 1990s, now see leading open source alternatives challenging these leaders. But open source's impact doesn't stop there. JBoss, the open source J2EE platform, is becoming the high-volume leader in the J2EE application space. More recently, Mono, the open source implementation of Microsoft .NET, promises to be a main strategic item for Novell. Mono may extend Microsoft's hegemony into Linux and open-source by countering J2EE's cross-platform strategic advantage.



BIO: Pierre Fricke, D.H. Brown's vice president of Application and Integration Infrastructure, extends the company's unique technical and strategic analysis into the J2EE, Microsoft .NET, and integration infrastructure space. After completing his M.B.A., Pierre became one of the leading strategists and marketing leaders in IBM focusing on interoperability, integration, WebSphere, Windows NT, UNIX, as well as Linux and open source.

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

VISIT www.sys-con.com/edge FOR TIMES & SCHEDULES

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

ces Conference & Expo

J-7 Migrating Enterprise Applications Between J2EE Application Servers

AJIT SAGAR, INFOSYS TECHNOLOGIES, LTD

Wednesday, February 16, 2005 5:00 P.M. – 5:50 P.M.

This session will provide guidelines, best practices, and a methodology to tackle a problem that is sapping the budgets of enterprise which have invested heavily in J2EE technology – the migration of enterprise applications between J2EE application servers. The strategy and planning for such initiatives is very complex and requires planning in advance. Enterprise applications, once deployed, have a multitude of dependencies, besides the dependency on Java APIs. The drivers for migration can include version upgrades, corporate agenda, maintenance costs, industry alliances, rapid upgrades to the J2EE platform APIs, etc. A planned migration ensures a successful implementation while minimizing the impact. This session will describe a strategy to plan for the migration of large portfolios of applications between application server vendors, between application server versions, and between hardware platforms. Real world examples of how this strategy has been applied in the industry will be provided.



BIO: Ajit Sagar is a senior technical architect with Infosys Technologies, Ltd., a global consulting and IT services company. Ajit has been working with Java since 1997, and has more than 15 years' experience in the IT industry. Ajit has served as *JDJ's* J2EE editor, was the founding editor of *XML-Journal*, and has been a frequent speaker at SYS-CON's Web Services Edge series of conferences. He has published more than 75 articles.

J-8 Design Patterns and Project Organizational Techniques for "Write Once, Debug Everywhere"

DR. MICHAEL JUNTAO YUAN, UNIVERSITY OF TEXAS

Thursday, February 17, 2005 9:00 A.M. – 9:50 A.M.

Unlike the wildly successful server-side Java technology, the true "write once run anywhere" vision has never been achieved for client-side Java. For Java developers offering end-to-end smart client-based SOA solutions, the development process is still "write once, debug everywhere." As the client-side platforms evolve from a handful of PC OSs to hundreds of devices with different OSs and UI characteristics, developers must leverage design patterns and innovative project organizational techniques to smooth the development and optimization process. This session will introduce more than a dozen of those design patterns and organizational techniques based on Michael's experience working with Nokia developers to develop scalable mobile enterprise applications.

BIO: Dr. Michael Juntao Yuan is an author, developer, and software architect for end-to-end mobile software. He is a contributing editor to *JDJ* and a frequent contributor to many developer forums and publications. He is the author of two books.

J-9 Using Grid Computing with Web Services and J2EE to Create Internet-based SOAs

KIERAN TAYLOR, AKAMAI

Thursday, February 17, 2005 10:00 A.M. – 10:50 A.M.

Service-oriented architecture, Web services, and J2EE technologies are dramatically changing the ways in which enterprises develop and deploy their Internet-facing applications. Because these applications potentially have a global user base, correctly architecting applications is a particular challenge. A combination of grid computing and utility computing offers a way to provide computing resources when and where they are needed, but developers must factor in certain considerations during design. This session will provide an in-depth overview of three real-world case studies using Grid computing in combination with Web services and J2EE to create Internet-based SOAs. It is designed for application architects and developers, and attendees will learn how applications can be designed to operate in a distributed computing environment such that performance and scalability problems are bypassed during deployment.



BIO: Kieran Taylor is currently the director of product management for Akamai Technologies and focuses strategy and direction for Akamai's emerging technologies, including its J2EE Internet applications delivery services. While at Akamai, Taylor has helped promote the open-standard Edge Side Includes (ESI), a markup language for dynamic content assembly and delivery at the edge, in use today by many companies.

J-10 Java Web Services Programming Tips & Tricks

ANDRÉ TOST, IBM

Thursday, February 17, 2005 3:00 P.M. – 3:50 P.M.

This session will present a collection of programming tips and tricks related to consuming and providing Web services in Java. This collection has been created by a number of developers and consultants and is the result of many real-life project experiences. We will focus on implementation aspects for Web services and not go into any detail on architecture or conceptual issues. In other words, these are the problems that developers face once they have started coding.



BIO: André Tost works as a solution architect in IBM's Software Group, where he focuses on SOA and Web services technology. He primarily engages with IBM's strategic business partners to enable them for J2EE and Web services. Before his current assignment, he spent many years in various IBM product development groups.

J-11 JCP Program: How the Java Technology Binary Software Standard is Managed and Evolves

ONNO KLUYT, SR. DIRECTOR & CHAIR, JCP PROGRAM, SUN MICROSYSTEMS

Thursday, February 17, 2005 4:00 P.M. – 4:50 P.M.

Why a Java technology standard? Why technology communities? This session will explore the circle of adoption and business opportunity from an IT Manager and IT developer perspective. How Java technology fits into these circles, and the significance of conformance and the "Write Once, Run Anywhere" promise. The role the Java Community Process (JCP) program plays by carefully focusing on binary compatibility and bringing together the community to agree on standards and the results of this effort – multiple implementations from many sources based on Java technology.



BIO: Onno Kluyt is chair of Java Community Process (JCP) Program, Senior Director, JCP Program and Jini Program, Sun Microsystems. Onno leads the Program Office for the JCP, which oversees the process, manages its membership, guides specification leads and experts through the process, leads the Executive Committee meetings, and manages the JCP.org Web site. Onno also heads up the JINI program, including JINI technology and JINI community.

Attend a **FREE*** Tutorial

Tuesday, February 15, 2005, 8 – 11 a.m.



AARON WILLIAMS

Developing Web Services Using Java Technologies

AARON WILLIAMS, MANAGER, JAVA COMMUNITY PROCESS PROGRAM OFFICE, SUN MICROSYSTEMS

Why do developers favor Java technology for developing Web services? Java technology is the most ready platform for Web services and service-oriented architectures, complete with interoperability, platform independence, and security built in. This tutorial will review several Java technologies for Web services standards that have been developed or are currently being updated through the Java Community Process (JCP) program as JSRs – Java API for XML-based RPC (JAX-RPC), Java API for XML Binding (JAXB), Java API for XML Processing (JAXP), Java API for XML Registries (JAXR), JSR 109, Implementing Java Web Services, JSR 172, J2ME Web Services and JSR 173, The Streaming API for XML (StAX). Also to be highlighted: JSR 244, Java 2 Platform, Enterprise, Edition 5.0 Specification, JSR 181, Web Services Metadata for the Java Platform, JSR 208, and Java Business Integration.



* Free Tutorials with VIP Preregistration ONLY!

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

VISIT www.sys-con.com/edge FOR TIMES & SCHEDULES

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

WEB SERVICES TRACK

WS-1 Ensuring Web Services Interoperability

CHRIS FERRIS, IBM

Tuesday, February 15, 2005 9:00 A.M. – 9:50 A.M.

Despite the open industry standards that underlie Web services, interoperability has been a key challenge for vendors and customers implementing Web services. One reason for this is that the relevant industry standards often permit multiple acceptable implementation alternatives. This presentation will discuss in detail the challenge of Web services interoperability and the role played by the premier industry organization formed to address it, the Web Services Interoperability Organization. In particular, the presentation will cover the critical importance of WS-I profiles to an organization's Web services initiatives, including the manner in which companies can put WS-I profiles immediately to work.



BIO: Chris Ferris is chair of the WS-I Basic Profile Working Group and a senior technical staff member with IBM's Emerging Technology Group. He has been actively engaged in open standards development for XML and Web services since 1999 and is an elected member of the OASIS Technical Advisory Board. Chris is also a coauthor and editor of the WS-Reliable Messaging specification.

WS-2 Web Services Standards: Going Behind the Mask

GLEN DANIELS, SONIC SOFTWARE

Tuesday, February 15, 2005 10:00 A.M. – 10:50 A.M.

Web services and service-oriented architectures (SOAs) are emerging as an integral part of the enterprise IT strategy. According to a recent IDC study, Web services – related revenue is expected to triple from \$1.1 billion worldwide in 2003, to \$3.4 billion in 2004, and \$16.6 billion by 2008. As SOAs proliferate and the number of Web services added to them increases, standards will play an increasingly significant role. This session will look at the state of key Web services standards such as WS-Choreography, WS-Reliability and WS-ReliableMessaging, SOAP/MTOM/XOP, WSDL, XPath, XQuery, and WS-Notification as well as related Java standards and open source efforts. It will also look at the organizational impact of standards adoption in the industry.



BIO: Glen Daniels is manager of standards and consortia at Sonic Software and coauthor of *Building Web Services with Java*. He has been working with Web services technologies since their inception in the late '90s, and in addition to developing products and helping to found Apache's Axis project, he has been an active participant in standards bodies such as the W3C, and a member of the SOAPBuilders interoperability group.

WS-2B Solving Complex Business Problems Through SOA

JOHN DALY, NETNUMINA

Tuesday, February 15, 2005 3:00 P.M. – 3:50 P.M.



BIO: John Daly is a recognized leader in service oriented architecture, enterprise wide integration and business intelligence. As Managing Director of netNumina, a leading firm focused on financial service, pharmaceutical and other fortune 100 companies, he is responsible for helping companies solve complex business problems by delivering sophisticated technology solutions that combine proven technologies with cutting edge thinking. John's ability to deliver technological solutions specifically designed for the user experience separates his approach from those of netNumina's competitors.

WS-3 The XML Data Challenge

JONATHAN BRUCE, DATADIRECT TECHNOLOGIES

Tuesday, February 15, 2005 4:00 P.M. – 4:50 P.M.

Most businesses store and query data with relational databases but need to use Extensible Markup Language (XML) to exchange and display data on the Web and with vendors and partners. As a result, programmers need to deal with both relational and XML data, often at the same time. Emerging standards such as XQuery, XQJ, and SQL/XML, promise to revolutionize data exchange and the ways applications are developed, deployed, and utilized. Learn the key facts about these standards, including what they mean, when

they will be available, and what you, the developer, can do to prepare.

BIO: As the XML Product Manager at DataDirect Technologies, Nancy Vodicka is responsible for DataDirect Connect for SQL/XML, a database-independent SQL/XML implementation that is currently shipping, and DataDirect XQuery, a database-independent XQuery implementation that is currently in development. Nancy has more than 15 years experience in the software industry working with technologies such as XML, Web services, relational databases, and SQL.

WS-4 How To Bulletproof Your Web Services

DAVID MCCAWS, PARASOFT

Wednesday, February 16, 2005 9:00 A.M. – 9:50 A.M.

Web services are gaining industry-wide acceptance and usage and are moving from proof-of-concept deployments to actual usage in mission-critical enterprise applications. Web services range from major services such as storage management and customer relationship management to much more limited services such as furnishing stock quotes or providing weather information. As companies and consumers begin to rely more and more on Web services, the need for developing reliable, high-quality Web services is even stronger. This session will explain issues specific to Web services and will illustrate solid engineering and testing practices required to ensure complete Web service functionality, interoperability, and security. Whether creating Web services from scratch or integrating a legacy back-end server via Web services, the practices and principles outlined in this session will be of great benefit.



BIO: David McCaw has over eight years of experience in helping software development teams improve quality throughout the development process. Over the last three years, he has led the Parasoft Web Services Solutions team, which has developed an industry-leading approach for Web services testing. He has implemented Web service quality solutions for development groups in organizations such as Sabre-Holdings, Yahoo! Overture, and McGraw-Hill. McCaw has an extensive background in the areas of Java and Web service reliability, performance, and security. He is involved with OASIS and WS-I, and is a frequent speaker at industry events.

WS-5 The Role of Policy in Web Services Integration – It's More Than Just Security

TOUFIC BOUBEZ, LAYER 7 TECHNOLOGIES

Wednesday, February 16, 2005 10:00 A.M. – 10:50 A.M.

Too often today the preferences, terms, and conditions describing how a Web service behaves when discovered and invoked is programmed right into the business logic. Hard-coding this behavior logic however introduces cost, complexity, and rigidity into a Web services architecture. A better approach is to abstract a Web services usage "policy" out of code where this metadata can be managed as need be. This session introduces the concept of Web Services Policy and describes how the construct can be used to implement a more customized and versatile Web service infrastructure.



BIO: Toufic Boubez is a well-respected and renowned Web services visionary. Prior to cofounding Layer 7 Technologies, Toufic was the chief Web services architect for IBM's Software Group and drove their early XML and Web services strategies. He is a sought-after presenter and has chaired many XML and Web services conferences. He is an author of many publications and his most recent book is the top-selling *Building Web Services with Java: Making Sense of XML, SOAP, WSDL, and UDDI*.

WS-6 B2B Policy Enforcement: The Third Rail of SOA Implementation

ALISTAIR FARQUHARSON, DIGITAL EVOLUTION

Wednesday, February 16, 2005 4:00 P.M. – 4:50 P.M.

One of the great benefits of a service-oriented architecture is the ability it gives you to extend programmatic, integration capabilities to business partners. Going beyond simple sharing of data with partners, SOA enables true B2B application integration. At the same time, this capability creates a vexing security policy enforcement dilemma. How can you be sure that a user from a partner organization is actually authorized to integrate with your applications? How can you authenticate that user? Do you even want that headache in the first place? This session will discuss the issues that arise in B2B security policy enforcement and explore several proven approaches

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

Services Conference & Expo

to solving the problem. In particular, it will focus on the emerging technology of XML Virtual Private Networks (XML-VPNs) and their potential to mitigate security policy enforcement issues in B2B SOA implementations.



BIO: Alistair Farquharson is the CTO of Digital Evolution, where he spearheads product development and provides thought leadership to enterprise customers implementing Web services. His skills span many industries and include designing and implementing system architectures, as well as spearheading initiatives such as development/team lead. He is an expert in custom-application development, distributed environments, architecting scalable hardware and software applications and systems, and Web services application development.

WS-7 Driving SOA Governance

BRENT CARLSON, LOGICLIBRARY

Wednesday, February 16, 2005 5:00 P.M. – 5:50 P.M.

In the past year, Web services and service-oriented architectures (SOAs) have become mainstream because of their ability to provide business agility and flexibility through integration, productivity, and reuse. With SOA enablement on the rise, IT groups must address SOA governance as a means of controlling what and how services located within an SOA are deployed. This session will discuss SOA governance, specifically how an organization can manage and control assets and artifacts located within an enterprise, while ensuring that deployed assets meet an organization's business and technical architectural standards. It will also outline governance best practices such as monitoring the UDDI publish process in order to seamlessly tie together the development and operational views of Web services within the enterprise.



BIO: Brent Carlson drives the development and delivery of LogicLibrary's products. He is a 17-year veteran of IBM, where he served as lead architect for the WebSphere Business Components project and held numerous leadership roles on the "IBM San Francisco Project." He is a member of the Eclipse Board of Stewards and a BEA Regional Director.

WS-8 SOA: From Pattern to Production

DAVID CHAPPELL, SONIC SOFTWARE

Thursday, February 17, 2005 9:00 A.M. – 9:50 A.M.

Service-oriented architecture (SOA) represents the opportunity to achieve broad-scale interoperability, while providing the flexibility required to continually adapt technology to business requirements. No small feat, particularly when one considers the extent and complexity of today's IT environments. As both a technology concept and IT discipline, the challenge inherent in SOAs is maintaining the right architectural approach. If all services in an SOA are treated as interdependent point-to-point interfaces, then the complexity of implementing and maintaining them in this spaghetti-like architecture becomes enormous. The enterprise service bus (ESB) has emerged as one of the first true SOA product offerings, bringing SOA from pattern to production. ESBs provide a framework for building and deploying an event-driven, enterprise SOA and accommodates the configuration, hosting, and management of integration components as services across the business.



BIO: VP and chief technology evangelist for Sonic Software, Dave Chappell has over 18 years of experience in the software industry covering a broad range of roles including R&D, code-slinger, sales, support, and marketing. He also has extensive experience in distributed computing, including message-oriented middleware, CORBA, COM, and Web application server infrastructure.

WS-9 High Performance Web Services – Tackling Scalability and Speed

SAMEER TYAGI, SUN MICROSYSTEMS

Thursday, February 17, 2005 10:00 A.M. – 10:50 A.M.

Web services facilitate application-to-application integration and interoperability across different platforms. However, critics usually point to an inefficient processing model and bandwidth requirements for developing Web services. This is often cited as a reason why Web services cannot perform and scale well in production environments. This session takes a detailed look at performance and scalability issues around Web services in the real world, as well as strategies that architects and developers

can adopt to mitigate such risks in these applications. Some analytical and modeling strategies that enable acceptable application performance will also be covered.



BIO: Sameer Tyagi works as a senior Java architect with Sun Microsystems. He remains focused on architecture, design, and implementation of large-scale enterprise applications with Java technology. His publications include industry periodicals and books on Java and J2EE technologies including *Java Web Services Architecture*.

WS-9B Effective Risk Abatement and Success in a Service-Oriented World

PAUL LIPTON, COMPUTER ASSOCIATES

Thursday, February 17, 2005 10:00 A.M. – 10:50 A.M.

IT leaders are hoping to leverage the benefits inherent in Web services and Service-Oriented Architectures (SOA) to enable their businesses to be far more competitive and to find new operational efficiencies. But, can we depend on these new technologies and approaches? Management and security are a common concern today and this session provides the necessary background and perspective on both the business and the technical issues. We will examine important principles and recommendations using real-world examples to illustrate key concepts.



BIO: Paul Lipton is a senior architect in the Web services and application management team at Computer Associates (CA) as well as a Strategist in the Office of the CTO. He has been an architect and developer of enterprise systems for over 20 years, and has worked closely with key CA customers to solve important business challenges through the creation of manageable, mission-critical distributed solutions.

WS-10 So You Want an SOA: Best Practices for Migrating Toward Service Orientation in the Enterprise

ERIC NEWCOMER, IONA

Thursday, February 17, 2005 3:00 P.M. – 3:50 P.M.

Replacing complex, monolithic applications with nimble applications built from exposed services promises increased developer productivity, greater flexibility, and ultimately reduced cost. The adoption of Web services and SOA can also remove a significant level of complexity and integration problems from enterprise application development projects. But, as with any large-scale project, IT departments must have the right plan and the right resources in place to ensure a successful transformation of their computing infrastructure. This article will explore what IT organizations need to know to be successful in their attempts to migrate the enterprise to a service-oriented architecture.



BIO: In the role of chief technology officer at IONA, Eric Newcomer is responsible for IONA's technology roadmap and the direction of IONA's e-business platforms as relates to standards adoption, architecture, and product design.

WS-11 Four Abilities SOA Will Lack Without a Registry

LUC CLEMENT, SYSTINET

Thursday, February 17, 2005 4:00 P.M. – 4:50 P.M.

A service-oriented architecture (SOA) is the design blueprint for seamless connectivity between business processes and IT infrastructure, enabling innovation and improving productivity. SOA provides the most efficient, standard way to dynamically interoperate with any customer, supplier, product or employee. SOA makes integration intrinsic. Web services are the foundation building blocks of an SOA, and they are already proliferating inside most enterprises. In an SOA, Web services become business services with the ability to perform a particular function or access data dynamically. This presentation will discuss the four abilities that a registry provides for an SOA.



BIO: Luc Clement is director of product marketing, SOA Registry for Systinet. He is also cochair for the UDDI Specification Technical Committee. Formerly Microsoft UDDI Program Manager, Luc is well known in the UDDI community and has been heavily involved with the UDDI specification for several years.

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

VISIT www.sys-con.com/edge FOR TIMES & SCHEDULES

The program, including topics and times, is subject to change. Please refer to www.sys-con.com for all updates.

International Web Services Conference & Expo



For Assistance Tel: 201-802-3066 Fax: 201-782-9601

CONFERENCE & EXPO: February 15 - 17, 2005

John B. Hynes Veteran Memorial Convention Center • Boston, MA

THREE WAYS TO REGISTER FOR CONFERENCE

- 1) **On the Web:** Credit Cards or "Bill Me" Please make checks payable to SYS-CON Events
- 2) **By Fax:** Credit Cards or "Bill Me" 201-782-9601
- 3) **By Mail:** 135 Chestnut Ridge Road, Montvale, New Jersey 07645, Attention: Registration

Please note: Registrations are not confirmed until payment is received.

Please complete sections 1, 2, 3 and 4

1 YOUR INFORMATION (Please Print) Mr. Ms. Mrs.

First Name _____ Last Name _____
 Title _____
 Company _____
 Street _____
 Mail Stop _____
 City _____
 State _____ Zip _____ Country _____
 Phone _____
 Fax _____ E-Mail _____

2 PAYMENT METHOD: (Payment in full due with registration)

Check or Money Order Enclosed (Registration confirmed upon receipt of payment)
 Check # _____ Amount of Check \$ _____
 Charge my Visa MasterCard American Express Discover
 Name on card _____
 Card # _____ Exp. Date _____
 Signature _____
 Billing Address (if different from mailing address) _____

3 PLEASE INDICATE YOUR CONFERENCE CHOICE:

Total Registration fee \$ _____

	Before 12/31/04	Before 1/21/05	Before 2/11/05	On Site
<input type="checkbox"/> 3D: Three Day Conference Good for all three days of the conference program, including keynotes, all conference sessions.	\$1,195.00	\$1,395.00	\$1,495.00	\$1,695.00
<input type="checkbox"/> 2D: Two Day Conference Good for two days of the conference program, including keynotes, all conference sessions.	\$995.00	\$1,195.00	\$1,295.00	\$1,395.00
<input type="checkbox"/> 1D: One Day Conference Good for one day of the conference program, including keynotes, all conference sessions.	\$595.00	\$695.00	\$695.00	\$795.00
<input type="checkbox"/> VIP PASS FREE Tutorials with VIP preregistration	FREE	FREE	FREE	\$75.00
<input type="checkbox"/> FREE - Novell Tutorial (Feb. 15)				
<input type="checkbox"/> FREE - Developing Web Services Using Java Technology (Feb. 15)				
<input type="checkbox"/> FREE - The Next Generation of Visual Studio (Feb. 16)				
<input type="checkbox"/> FREE - Security, The New Reality (Feb. 17)				
<input type="checkbox"/> EO: Expo Hall Only Includes keynotes, expo presentations, and workshops	FREE	FREE	FREE	\$50.00

SPECIAL GROUP RATES

Registering a group of four or more? Please call Jim Hanchrow at 201-802-3066, and receive an additional discount.

4

A. Your Job Title

- Senior Software Developer/Engineer/Architect/Analyst
- Director of IT/Development/Engineering
- Software Developer/Engineer/Architect
- Systems Architect/Engineer/Analyst
- Software Consultant/Analyst
- Project Manager/Team Leader
- Business Development
- General Manager/Division or Department Head
- CTO,CIO, Chief Architect
- CEO, President, Owner, Partner/Principal
- VP/Technical Director
- Network Administration/Manager/Consultant
- Web Developer/Programmer
- Database Administrator
- VP/Sales & Marketing Director /Product Management
- Other (please specify) _____

B. Business/Industry

- Agriculture/Mining/Oil/Gas
- Computer Hardware and Electronics
- Construction/Architecture/Engineering
- Consulting (non-computer)
- Education
- Banking/Insurance/Accounting
- Government (including Military)
- Healthcare/Pharmaceuticals/Biotech/Biomedical
- Internet/Web/e-Commerce
- Manufacturing and Process (other than computer related)
- Media/Marketing/Advertising
- Non-Profit/Trade Association
- Real Estate/Legal
- Telecommunications
- Transportation/Rail/Sea/Air/Ground
- Travel/Hospitality/Entertainment
- Utilities
- Wholesale/Trade/Distribution/Retail
- Other (please specify) _____

C. Total Number of Employees at Your Location and Entire Organization (check all users that apply):

	Location	Company
10,000 or more	01 <input type="checkbox"/>	01 <input type="checkbox"/>
5,000 - 9,999	02 <input type="checkbox"/>	02 <input type="checkbox"/>
1,000 - 4,999	03 <input type="checkbox"/>	03 <input type="checkbox"/>
500 - 999	04 <input type="checkbox"/>	04 <input type="checkbox"/>
100-499	05 <input type="checkbox"/>	05 <input type="checkbox"/>
100 or less	06 <input type="checkbox"/>	06 <input type="checkbox"/>

D. Please indicate the value of software products and services that you recommend, buy, specify or approve over the course of one year:

- \$10 million or more
- \$1 million - \$9.9 million
- \$500,000 - \$999,999
- \$100,000 - \$499,999
- \$10,000 - \$99,999
- Less than \$10,000
- Don't know

- \$10 billion or more
- \$1 billion - \$9.9 billion
- \$100 million - \$999 million
- \$10 million - \$99.9 million
- \$1 million - \$9.9 million
- Less than \$1 million
- Don't know

E. What is your company's gross annual revenue?

F. Do you recommend, specify, evaluate, approve or purchase mainframe products or services for your organization?
 01 Yes 02 No

G. Which of the following products, services, and/or technologies do you currently approve, specify or recommend the purchase of?

- Application Servers
- Web Servers
- Server Side Hardware
- Client Side Hardware
- Wireless Device Hardware
- Databases
- Java IDEs
- Class Libraries
- Software Testing Tools
- Web Testing Tools
- Modeling Tools
- Team Development Tools
- Installation Tools
- Frameworks
- Database Access Tools / JDBC Devices
- Application Integration Tools
- Enterprise Development Tool Suites
- Messaging Tools
- Reporting Tools
- Debugging Tools
- Virtual Machines
- Wireless Development Tools
- XML Tools
- Web Services Development Toolkits
- Professional Training Services
- Other [Please Specify] _____

SYS-CON Events, Inc. and SYS-Con Media make no warranties regarding content, speakers or attendance. The opinions of speakers, exhibitors and sponsors do not reflect the opinion of SYS-CON Events and SYS-CON Media and no endorsement of speakers, exhibiting companies products, or sponsors is implied.

Speakers, sessions and schedule are subject to change without prior notice.

No solicitation by anyone other than official exhibitors, sponsors or marketing partners is permitted. Such behavior is cause for expulsion without refund.



If you require special assistance covered under the Americans with Disabilities Act, please call 201-802-3066 by February 4, 2005.

CANCELLATIONS, SUBSTITUTIONS, REFUNDS

Fax written request to SYS-CON Registration 201-782-9601. Requests for refunds received prior to January 7, 2005 will be honored, less a 10% handling charge; requests received after January 7, 2005, and before February 4,

2005, will be honored less a 20% handling charge. No requests for refunds will be honored after February 4, 2005. Requests for substitutions must be made in writing prior to February 11, 2005. No one under 18 is permitted to attend. No warranties are made regarding the content of sessions or materials.

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

Some know
exactly where technology is headed. Others are there before it

arrives.



The Intel Developer Forum is where innovators kick the tires on new ideas, build new relationships and unleash their vision of where it's all heading. See Intel's newest platforms emerging. Watch new communities organizing. Experience new ecosystems forming.

**Tomorrow's Platforms:
Your Possibilities Unleashed.**

San Francisco March 1-3, 2005

This is the IDF you won't want to miss.

REGISTER NOW FOR VALUE PRICING:

www.intelidf.com/reg4

1-888-261-0971

Priority Code: PRT1ALL



Onno Kluyt

From Within the Java Community Process Program

Java technology interfaces for Web services

Welcome to the February edition of the JCP program column. Each month you can read about the Java Community Process program: newly submitted JSRs, new draft specs, Java APIs that were finalized, and other news from the JCP. JCP members have been busy since my last column, submitting over six new JSRs! This time around I'll cover some new JSRs recently submitted in the Web services area, and some JSRs from the telecommunications community and from the Micro Edition folks.

Web Tier

In the past month or so four new JSRs have been submitted to the JCP program that focus on Java technology interfaces for Web services. You can read more about these JSRs in my Web services column also in this issue.

JSR 267, a JSP tag library for Web services, proposes a happy union of standard Web services protocols and JavaServer Pages (JSP) that will help JSP page authors integrate Web services into their Web applications with minimal effort and quick turnaround times.

JSR 265, API for Utilizing Web Services Policy, aims to standardize a basic framework in Java for utilizing the Web services constraints and capabilities.

JSR 262, Web Services Connector for Java Management Extensions (JMX) Agents, defines a connector for the JMX Remote API that uses Web services to make JMX instrumentation available remotely.

JSR 261, Java APIs for XML-Based Web Services Addressing (JAX-WSA) 1.0 specification, will define APIs and a framework for supporting transport-neutral addressing of Web services.

In addition, two JSRs that directly support the ease of development goals of J2EE 5.0 (JSR 244) are JavaServer

Faces 1.2 (JSR 252) and JSP 2.1 (JSR 245) technologies. These two JSRs recently completed the Early Draft Review in the JCP program and are part of the J2EE 5.0 Umbrella JSR. The focus of these JSRs is to improve the alignment of the J2EE Web tier and to further enhance ease of development. A cornerstone of this is a new, unified expression language API that both technologies share.

You can find out more about Java technologies being developed for Web services at the Web Services Edge 2005 East Conference. The JCP is an association and corporate sponsor of the event and we are holding our quarterly JCP training and face-to-face Executive Committee (EC) meeting nearby, hosted by IBM. We will have a booth in the Expo Hall (#405) and a free tutorial on developing Web services with Java technologies on Tuesday, from 8 a.m.–11 a.m. I will also be leading a panel session with JCP EC members on Tuesday afternoon at 3 p.m. and a technical session on Thursday afternoon at 4 p.m. Hope to see you there!

Telecommunications Community – OSSJ and JAIN

Two new JSRs have been submitted and are supported by the OSS through the Java community. JSR 264, Order Management API, defines an interface to an Order Management component, with specific extensions for activation and work order management use cases. JSR 263, Fault Management API, provides for the monitoring and management of alarms within an information or a telecommunications network. Both JSRs were approved by the JCP EC in January.

In addition, the JAIN Community has two JSRs completing Public Review in February. JSR 186, JAIN Presence, is a generic and protocol-agnostic API

for Presence that provides a standard portable and secure interface to control, manage, and manipulate Presence information between Presence clients and servers. JSR 187, JAIN Instant Messaging, a protocol-agnostic API for Instant Messaging, provides a standard portable and secure interface to control, manage, and manipulate instant messages between clients through the use of Presence servers.

The expert group for JSR 79, JAIN MEGACO API, has completed its work and published the Final Release of the specification. This JSR specifies the JAIN API through which access to protocol stacks that support the MEGACO/H.248 protocol may be achieved.

On the J2ME Side

For the J2ME environment, one new JSR has been submitted and approved by the EC – JSR 266, Unified Message Box Access API (UMBA-API). This JSR will define an API to access and manage the mobile device's message boxes and their contents.

The Expert Groups of two J2ME JSRs have recently posted Propose Final Drafts. JSR 238, Mobile Internationalization API, defines an API that provides culturally correct data formatting, sorting of text strings, and application resource processing for J2ME MIDlets running in MIDP over CLDC. JSR 211, Content Handler API, will enable J2ME applications to handle multimedia and Web content, and can give developers and users a seamless and integrated user environment on mobile phones and wireless devices.

That's it for this month. I'm very interested in your feedback. Please e-mail me with your comments, suggestions, and questions. I look forward to seeing you at the Web Services Edge 2005 East Conference! ☺

Onno Kluyt is the director of the JCP Program Management Office, Sun Microsystems.
onno@jcp.org

ENTERPRISE CHARTING SOLUTIONS

Need More Than A Pretty Face?



For Serious Enterprise-Level Data Visualization & Analysis

Multi-Platform
Ease of Use
Outstanding Support
...and A Pretty Face!

Extensibility
Performance
Scalability



Chart FX

www.softwarefx.com

You didn't compromise on your infrastructure... show your applications the same respect.



RICH WEB APPS

Need desktop interactivity within a web browser?



LIVE DATA

Wish your browser-based apps had Excel-like functionality?



UNIVERSAL DEPLOYMENT

Imagine building an application once, deploying it anywhere?



INDUSTRIAL STRENGTH

Want enterprise-scale Web apps that are reliable, secure and support 1,000's of concurrent users?

Nexaweb's rich client platform makes your apps available anytime, anywhere.

We know you're tired of being asked to do the impossible. We can help.

Nexaweb lets you leverage all the work you've done in client/server and deploy your apps on the Internet, without compromise.



“Nexaweb was able to help us deliver significant cost savings while enhancing overall system performance over infrastructure alternatives.”

Mike Kistner, CIO and
Vice President, Best Western