WILL OPEN AJAX CLIMB TO THE TOP? PG 50

WebSphere® JOURNAL

www.WebSphere.SYS-CON.com

The World's Leading Independent WebSphere Resource

FEBRUARY 2006 VOLUME 5 ISSUE 2

The Six Principles of Systems Engineering

IBM's Rational Rules Developed Over 10 Years

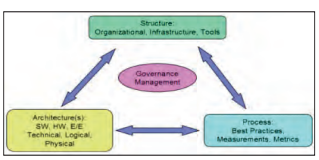

BY MURRAY **CANTOR**
GENE **ROOSE**

PAGE 6

Integrate EJB Services

W ith WebSphere Process Server

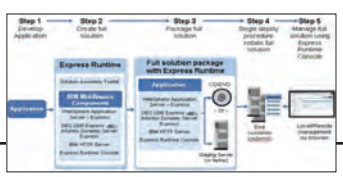

BY PETER **XU**

PAGE 26

Under the Hood of IBM Workplace Collaboration Services

Many Layers Mean Rich Functionality

BY BOB **BALABAN**

PAGE 30

Demystifying Class Loading Problems

Part One of a Four-Part Article

BY LAKSHMI **SHANKAR**
SIMON **BURNS**

PAGE 38

PRESORTED
STANDARD
US POSTAGE
PAID
ST. CROIX PRESS

NEW!

CodePro AnalytiX™

Continuous Collaborative Code Analysis for Rational®, WebSphere® & Eclipse

Imagine your software team creating higher quality, more reliable, Java systems... with less time and money.

Sound like a fantasy?

...*CodePro AnalytiX* lets you do exactly that, or

Your money back.

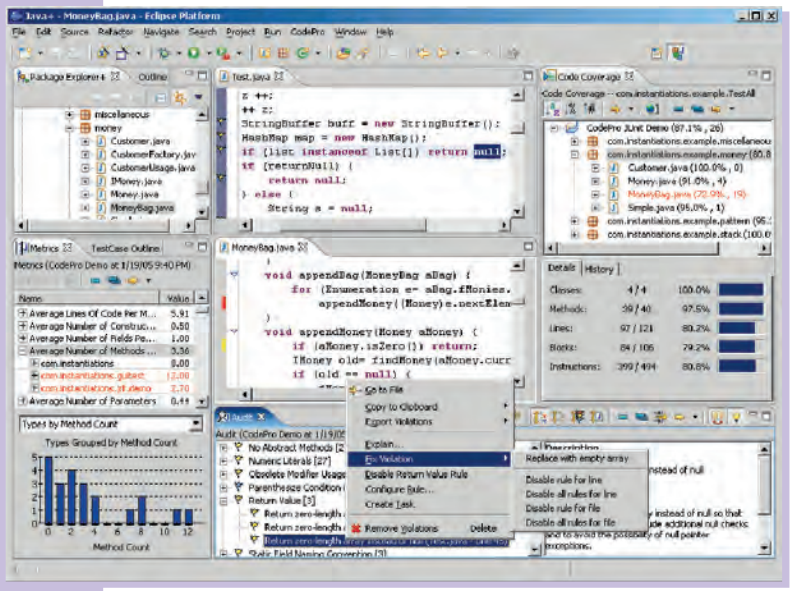

***AnalytiX* automates the process of building quality into your software as it's developed. Using the latest software development technology *AnalytiX* monitors Java code development in real-time, as its written. It identifies and automatically fixes hundreds of issues that would reduce quality and reliability.**

Powerful wizards generate unit and regression tests and analyze code coverage, helping drive out problems early in the development process.

Less experienced developers code like experts, experts create tight, maintainable systems.

***AnalytiX* also allows you to define, distribute, and enforce corporate coding standards and quality measures across your development teams...no matter where they're located.**

Early identification and correction of code quality issues dramatically improves your application development process and reduces the time and money it takes to develop your high performance Java systems.

From the development tool experts at Instantiations. No one outside of IBM has longer or deeper experience in creating Eclipse-based Java development tools.

Put our experience to work for your team.

www.instantiations.com

1-800-808-3737

Key Product Capabilities

- **Detect and correct code quality issues ...*Automatically***
- **Define, distribute and enforce code quality standards across development teams**
- **700+ rules and metrics/300+ Quick Fixes**
- **Code metrics with drilldown and triggers**
- **Custom rules and rule sets**
- **JUnit test case generation**
- **Code coverage analysis**
- **Dependency analysis and reporting**
- **Integrated team collaboration**
- **Javadoc analysis and repair**
- **Seamless integration with IBM Rational Developer, WebSphere Studio and Eclipse**

Risk Free Trial Copy

www.instantiations.com/analytix

eclipse READY. IBM Business Partner eclipse FOUNDATION MEMBER

instantiations

Open Up!

Portals–Get what you want easily

Ask the Experts

Prolifics is 1 of only 3 WebSphere Service Providers retained by IBM. Open up with portals! With 26 years of industry expertise, Prolifics' certified WebSphere portal experts deliver a personalized, secure, unified view to your enterprise assets.

Our portal solutions build agility into your business — enabling you to reach new markets, achieve customer satisfaction beyond expectation, and ease collaboration within your enterprise. Our teams of specialists work together with your business experts from gathering the business requirements, translating them to a technical specification, to delivering the final production application.

Portal Executive Assessment

Prolifics experts help you articulate the value of portals to your business sponsors and perform a quick business review identifying portal benefits as it applies to your organization.

Portal Workshops and Proof of Concepts

Prolifics demonstrates valuable portal features and translates the breadth of your portal objectives into a step-by-step roadmap to meet your current and future needs.

Portal Jump Start

Prolifics' consultants work with you from the start to install, configure and jump start your team on using WebSphere Portal Software.

Portal Development and Mentoring

Prolifics' innovative approach to building portals drives down your overall cost and transforms the way you conduct business. We mentor and arm your team with the skills to manage, maintain and extend the portal environment.

Prolifics

Winner of 2003 IBM BUSINESS PARTNER LEADERSHIP AWARD

116 John Street, New York, New York 10038
800.675.5419
solutions4websphere@prolifics.com
www.prolifics.com

Prolifics® IBM Business Partner

Web 2.0 and the Future

BY ROGER STRUKHOFF

The new year is proceeding with what appears to be a renaissance or perhaps mini-bubble in the global IT industry. Use of the technically imprecise but no doubt seductive term "Web 2.0" is stimulating companies to think through what they need to deploy for and through their websites in ways unimaginable in the dark days of a few years ago.

Within this Web 2.0 context lies the emerging reality of open source software: its proponents, its political aspects, its communities, and most of all, its myriad downloads in myriad versions.

Although open source "threatens" to take away as much as 20 percent of the market within a few years, according to numerous sources, the fact is that open source is a mere reflection of an increasingly vibrant industry. Whether out of fear, expediency, or leadership, IBM continues to dance on the open source bandwagon with both feet, its most recent effort being what appeared to be a hasty assemblage of companies under the "Open AJAX" coalition, thus confirming its acknowledgement that open source is here to stay.

And is open source really the threat that many say it is? Probably not. Anecdotal evidence that I've collected suggests that many open source developers are independents, albeit often working on contracts for companies, who would otherwise be unemployed independents who cannot or refuse to buy high-priced software development tools. These independents are downloading open source software instead of buying nothing. Their market share is not being stolen from the proprietary environments.

Yet, once they get into companies and start to design and deploy applications, they end up influencing the purchase of numerous other IT goods and services, bringing a net gain to the overall market. And often, they end up interacting with, i.e., their projects end up interoperating with, many complex applications and infrastructures that are decidedly high-budget, thereby adding another net gain to the overall industry.

IBM, by positioning itself squarely in the open source arena while simultaneously dominating the traditional web services market through WebSphere, is one of the few companies that can have it both ways, and serve the entire spectrum of application development and administration. As one recent software executive pointed out to me, "you know, after all this, 70 percent of all the data in the world is still on mainframes."

There are still truly mission-critical issues and applications that cannot be cobbled together in any sort of loose way, but rather require the highest levels of tight integration and security. At the same time, there are terabytes of data on these big systems that need to be pushed outward. Squaring this circle is the crux of the challenge.

The IT world, as with the world in general, is a different place than it was in the late '90s, and a seemingly more complex one. Who would have thought in 1999 that we would one day be yearning for those simpler times! Yet just as the world is more complex, so are the opportunities more subtle. Web 2.0 offers hope for those who can grok the difference between publishing and syndication, between push and on-demand, between point-to-point and end-to-end, between comprehensive transmitability and asynchronous periodicity.

This issue of *WebSphere Journal* can hardly hope to addresss all these overriding issues (or are they meta-issues), but can take peeks at some of the facets of application development among IBM customers today.

Next month will see the completion of some of the thoughts started in this issue, as well as the commencement of new thoughts. The complexity seems almost Mandelbrotian, but the good news is that there appears to be infinite ways to solve the infinite challenges posed by enterprise IT problems (and opportunities) today.

ABOUT THE AUTHOR

Roger Strukhoff, editor-in-chief of *WebSphere Journal*, is West Coast Bureau Chief for the SYS-CON News Desk, and President of www.wdva.com. He spent 15 years with Miller Freeman Publications and The International Data Group (IDG), then co-founded CoverOne Media, a custom publishing agency that he sold in 2004. His work has won awards from the American Business Media, Western Press Association, Illinois Press Association, and the Magazine Publishers' Association. Read his blog at http://www.rssblog.linuxworld.com.

roger@sys-con.com

WebSphere® JOURNAL

ADVISORY BOARD

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

EDITOR-IN-CHIEF ROGER STRUKHOFF
CONTRIBUTING EDITOR PATTI MARTIN
PRODUCT REVIEW EDITOR JAY JOHNSON
EDITOR NANCY VALENTINE
ASSOCIATE EDITOR SETA PAPAZIAN

SUBSCRIPTIONS

For subscriptions and requests for bulk orders, please send your letters to Subscription Department.
888-303-5282 201-802-3012
SUBSCRIBE@SYS-CON.COM
Cover Price: $8.99/Issue Domestic: $149/YR (12 Issues)
Canada/Mexico: $169/YR Overseas: $179/YR
(U.S. Banks or Money Orders)
Back issues: $15 U.S. $20 All others

PRESIDENT AND CEO FUAT KIRCAALI
PRESIDENT, SYS-CON EVENTS GRISHA DAVIDA
GROUP PUBLISHER JEREMY GEELAN
SENIOR VP, SALES & MARKETING CARMEN GONZALEZ
VP, INFORMATION SYSTEMS ROBERT DIAMOND
VP, SALES & MARKETING MILES SILVERMAN
PRODUCTION CONSULTANT JIM MORGAN
FINANCIAL ANALYST JOAN LAROSE
CREDITS & ACCOUNTS RECEIVABLE GAIL NAPLES
ACCOUNTS PAYABLE BETTY WHITE
ADVERTISING DIRECTOR ROBYN FORMA
ADVERTISING MARKETING DIRECTOR DENNIS LEAVEY
ADVERTSING SALES MANAGER MEGAN MUSSA
ADVERTISING SALES MANAGERS DOROTHY GIL
KERRY MEALIA
LEAD DESIGNER ABRAHAM ADDO
ART DIRECTOR ALEX BOTERO
ASSOCIATE ART DIRECTORS ANDREA BODEN
TAMI BEATTY
LOUIS F. CUFFARI
VIDEO PRODUCTION FRANK MORICCO
WEB DESIGNERS STEPHEN KILMURRAY
VINCENT SANTAITI
SHAWN SLANEY
ONLINE EDITOR ROGER STRUKHOFF
CIRCULATION SERVICE COORDINATORS EDNA EARLE RUSSELL
LINDA LIPTON
CUSTOMER RELATIONS MANAGER BRUNI STAROPOLI

EDITORIAL OFFICES

SYS-CON Publications, Inc.
135 Chestnut Ridge Road, Montvale, NJ 07645
Telephone: 201 802-3000 Fax: 201 782-9637
SUBSCRIBE@SYS-CON.COM
WebSphere® Journal (ISSN# 1535-5914) is published monthly (12 times a year).
Postmaster send address changes to:
WebSphere Journal, SYS-CON Publications, Inc.
135 Chestnut Ridge Road, Montvale, NJ 07645

WORLDWIDE NEWSSTAND DISTRIBUTION
CURTIS CIRCULATION COMPANY, NEW MILFORD, NJ

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

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

SYS-CON MEDIA

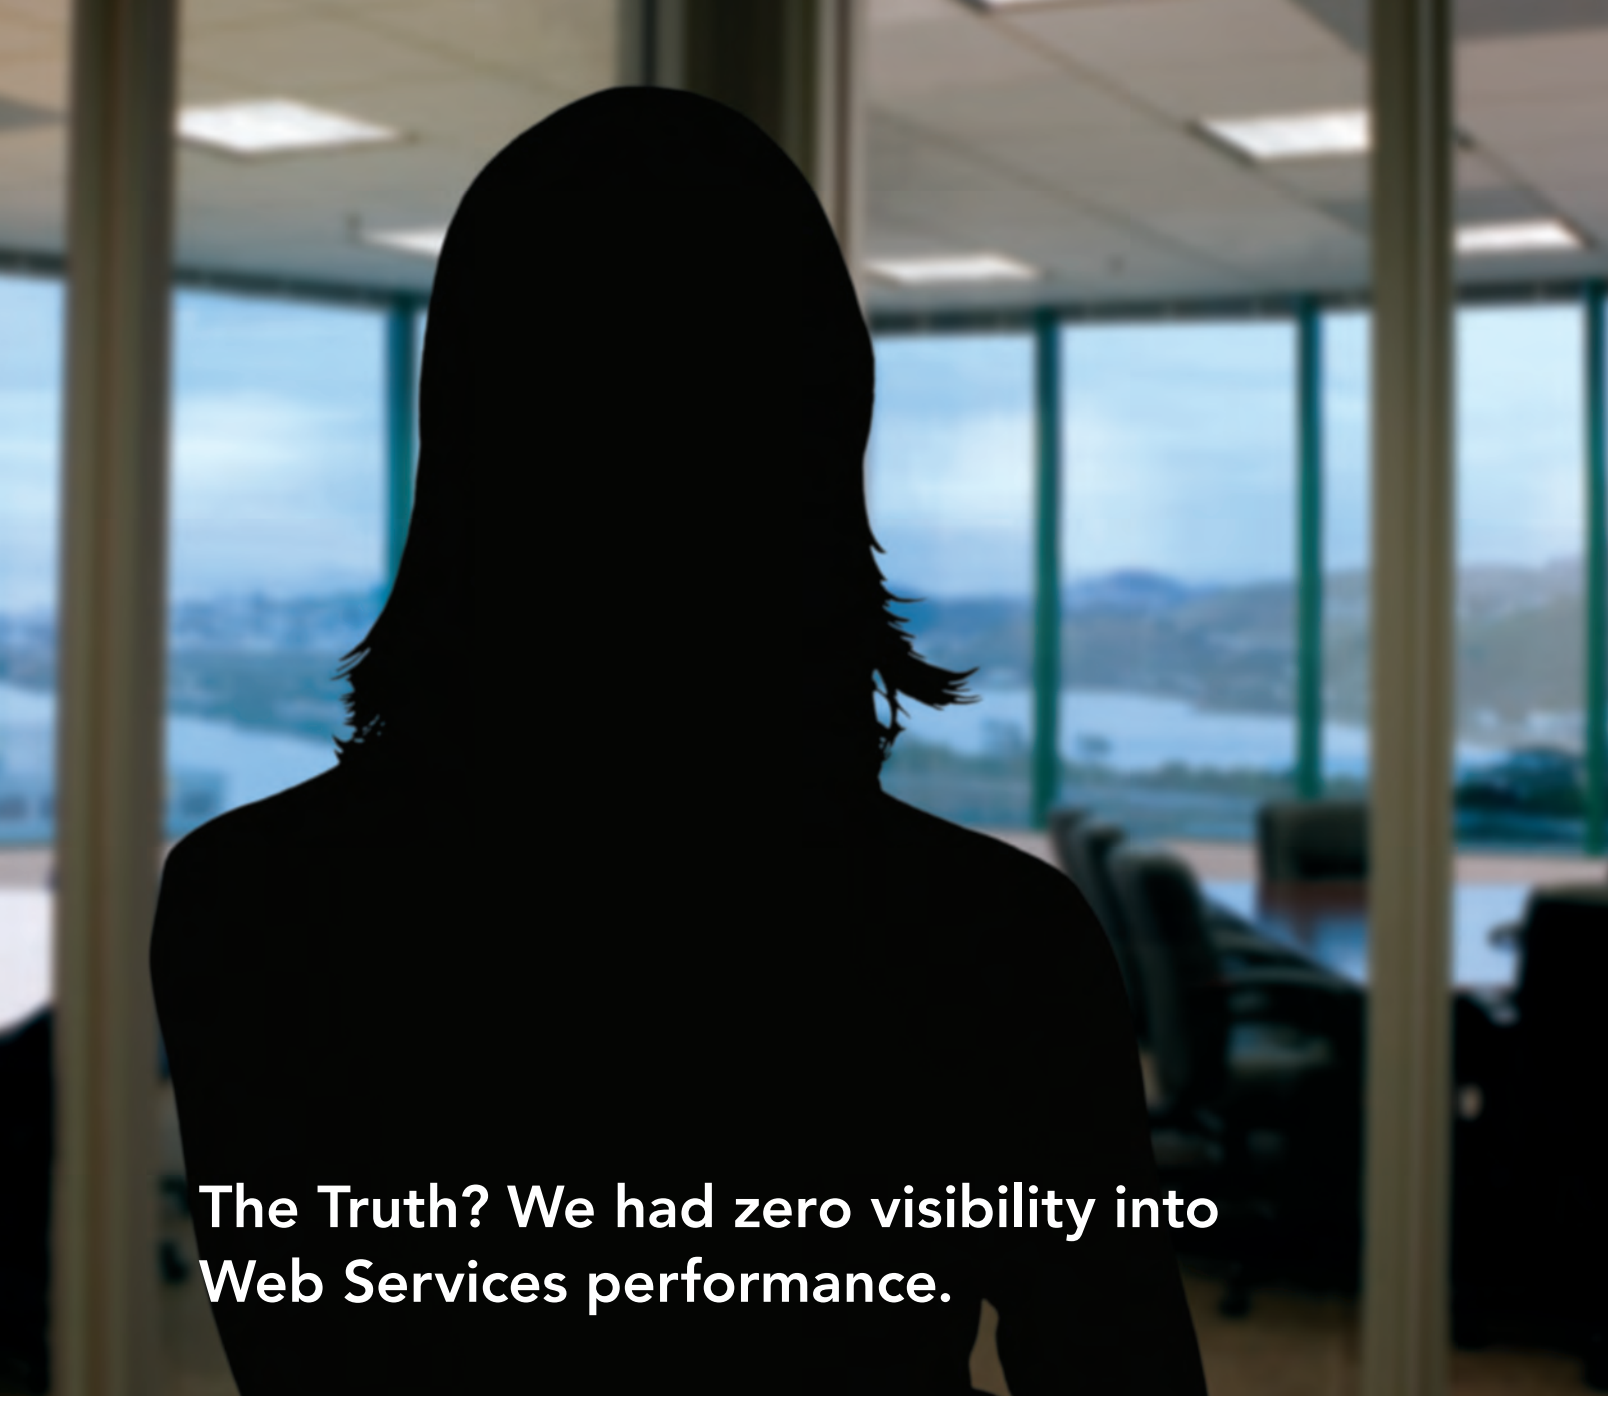

Wily can handle the truth. We've helped hundreds of organizations around the world manage critical Web Services and proactively identify problems before customers are affected. Using a single solution, you gain control over both your application infrastructure and the customer experience. To learn more about how Wily can help you achieve customer success with Web Services, visit **truth.wilytech.com**.

Get Wily.™

Enterprise Application Management

wily technology™

IBM's Rational Rules Developed Over 10-Year Period

The Six Principles of Systems Engineering

BY MURRAY CANTOR AND GENE ROOSE

IBM Rational's six principles of systems engineering are a set of high-level systems development guidelines derived from the careful analysis of successful, complex systems development engagements over the past ten years.

ABOUT THE AUTHORS

Murray Cantor is a member of the IBM Rational CTO team, taking the lead for Governance and Systems. Cantor formerly lead the IBM Rational field services group, promoting and extending Rational best practices, and working closely with customers on innovative ways to build and deliver systems more efficiently.

Gene Roose is a senior systems engineering consultant for IBM Rational, concentrating on model-driven systems development methods. He assists clients with solution analysis and design, architectural derivation and validation, and overall project management. His collaboration with Murray Cantor on a significant systems engineering process study for a major automotive concern provided foundational material for this article.

Although they are neither comprehensive nor mutually exclusive, they serve to highlight key areas of focus for organizations interested in quickly building expertise in complex systems development. They also serve as a "measuring stick" for assessing potential problem areas and the root causes underlying symptomatic project deficiencies or failures.

It is well-known and accepted that function, schedule, and cost are three key and mutually dependent aspects of project management -- make a change to one, and the effects often ripple through the other two.

A similar relationship exists in product and program management for complex systems development. As shown in Figure 1, the three key aspects are:

1. Systems architecture
2. Organizational structure, including the systems development infrastructure
3. Process, including workflows, best practices, and the like

These aspects collaborate to help achieve the following:

- "Build the *right* system and build the system *right*"

  Architecture focuses on building the right system and model-driven use-case based architectural decomposition lays the foundation for building the system right. Why is architecture singled out? Because it provides the foundation on which the system is ultimately constructed. The choices made in selecting the system's architecture become a lasting legacy – either good or bad.

  Here's a transportation example: The basic architecture used for centuries in building roads or other transportation elements (like rail tracks) was basically point-to-point (the shortest distance between two points can be considered optimal since it minimizes the road length and one would assume maximizes time to travel efficiency). Simple enough, but there was always at least a physical constraint: the terrain, which forced the path to avoid obstacles. There were also other potential constraints such as "hostile" territory or political boundaries.

  In any case, roads were built and paths crossed somewhat haphazardly. Move forward to the 20th century. Population growth generated additional constraints – namely street lights or traffic signs to optimize the flow of traffic through intersections. These are somewhat problematic on roads architected in the point-to-point avoid obstacle approach, especially where more than two roads intersect. A clearly "better" architecture has proven to be laying out roads in rectangular grids. This architecture and the many implementations are much easier to control from a signaling and traffic flow perspective.

  So build the right system and build it right – making it robust (scalable, extendible, easy to maintain, easy to use, and so on).

- "Do the right things, and do the things right"

  The systems development framework focuses on defining the right things to do throughout the product lifecycle to optimize business value and return on investment and best practices, processes/supporting tools and governance (requirements, quality, change, configuration, and program/project management) focus on ensuring things are done right.

MAKE ANSWERS TO PERFORMANCE PROBLEMS COME TO YOU.

OPNET Panorama offers powerful analytics for rapid troubleshooting of complex J2EE/.NET applications. Panorama quickly identifies how application, web, and database servers are impacting end-to-end performance. With Panorama, you can pinpoint the source of a problem, so time and money aren't spent in the wrong places.

The most successful organizations in the world rely on OPNET's advanced analytics for networks, servers, and applications.

OPNET Panorama™
Real-Time Application Analytics

www.opnet.com/pinpoint

OPNET®
Making Networks and Applications Perform™

OPNET Technologies, Inc. 7255 Woodmont Avenue, Bethesda, Maryland 20814 phone: (240) 497-3000 • e-mail: info@opnet.com • NASDAQ: OPNT

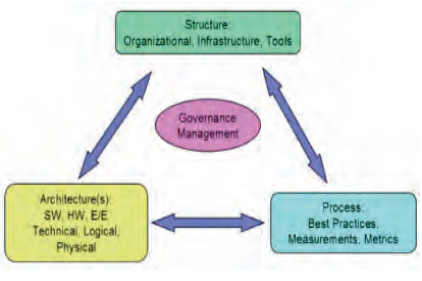

Figure 1: Three interdependent aspects of managing complex systems development

- "Apply the right resources, and apply the resources right"

"Doing" and "building" are action verbs; they require resources to perform the "doing" and the "building". The most significant enterprise resource is people, followed by supporting infrastructure. Acquiring, motivating, and retaining skilled personnel applies the right resources, and aligning the organization appropriately to optimize communication, collaboration, and effort applies the resources optimally. A sound enterprise infrastructure (including IT equipment and software) is one key aspect in ensuring the resources are applied optimally, once again maximizing business value and return on investment

The six principles of systems engineering address all three aspects described above. The three technical principles (noted below) focus on architecture and the derivation of system models, while the remaining principles provide the complementary infrastructures and workflows needed to optimize the technical development environment.

The six systems engineering development principles are:

1. Decompose systems, not requirements (technical).
2. Enable both separation and integration of «key systems development» concerns (technical).
3. Specifications flow up and down the architecture (technical).
4. Systems and components collaborate; so should development teams.
5. Development organizations should reflect product architectures.
6. Base the «development» lifecycle on removing risk and adding value.

Let's examine each of these principles, initially from a generic systems engineering view, and then specifically in the context of joint hardware/software development.

Decompose systems, not requirements

Since systems and software engineering principles and methods have been written about, discussed, and applied in countless cases, one might think that development teams have a common understanding of such terms as "system" and "systems engineering." Unfortunately, this is not the case; neither across enterprises within a specific industry nor across the product and functional domains within an engineering development community in a single enterprise.

According to INCOSE (the International Council on Systems Engineering), *systems engineering* is[2]:

An interdisciplinary approach and means to enable the realization of successful systems. It focuses on defining customer needs and required functionality early in the development cycle, documenting requirements, and then proceeding with design synthesis and system validation while considering the complete problem:

- Operations
- Performance
- Test
- Manufacturing
- Cost and Schedule
- Training and Support
- Disposal

A *system* is a group of interacting, interrelated, or interdependent elements forming a complex whole, which provides a set of services that are used by an enterprise to carry out a business purpose (mission). System components consist of hardware, software, data, and workers[3]. Put simply, a system is a complex entity that provides some tangible result of value. Systems engineering is a disciplined approach that helps us examine desired results and determine what can satisfy them. It also helps us determine how to do this within a set of business-critical constraints (cost, schedule, testing parameters, ease of manufacturing, and so on). The "what" is the *system*, the "how" is the *functional requirements*, the "business-critical constraints" are the *supplementary requirements*, and the "desired results" are products of *use cases*.

At the highest level of abstraction, we can view the system as a single entity (complex whole), or *black box*. However, to build a system, we must also be able to open the box, look inside, and see what set of interacting, interrelated, or interdependent elements make up the system. This is called the *white box* view. The primary task of systems (and software) engineering is to recursively move from black box to white box views, opening the system box, subsystem boxes, and sub-subsystem boxes until we ultimately determine the set of interacting and interdependent elements that can provide the system's desired results.

A close look at the INCOSE definition of systems engineering does not reveal a development methodology preference for realizing systems. Traditional requirements-driven methods, with their typically rigid and rigorously defined project tasks, milestones, resources, and schedules, seem to apply. However, as noted earlier, these methods were not designed with a focus on early risk mitigation, fluid requirements, or building systems for which the solution was initially indeterminate; or even at the limits of scientific or technical knowledge.

As we noted earlier, focusing primarily on the requirements and their functional allocation leads to early (and binding) decisions about technology and architecture, even if the architecture is not explicitly addressed. These early decisions can lead to undesirable results:

- Delivery of systems that fail to meet key stakeholder needs.
- Significant rework to address changing or misinterpreted requirements.
- Expensive "heroics" in the final development stage (systems integration) to "rescue" runaway projects.
- Worst of all, scrapping the system entirely, after heroic efforts fail.

Imagine seeing EAI's impact on business process performance in real-time...

Now imagine implementing it in minutes, *instead of* months

To get your free copy of the informative new article

"3 Steps to Uncomplicating Business Process Management"

go to www.nastel.com/uncomplicate
or call 800-580-2344 or
+631-761-9100

Business happens in real time.
If something impacts your business processes, you need to know. *Now.*

AutoPilot™ from Nastel monitors and manages all of the complex events that comprise your business-critical transactions, giving you deeply granular visibility from end to end for all your business processes. AutoPilot lets you define customized, actionable alerts in the context of each user's specific business function to streamline your business processes and eliminate costly failures, bottlenecks and redundancies. With AutoPilot, you can see the impact of EAI on your business process performance in real time, on-demand.

Unlike other monitoring software, though, AutoPilot's unique architecture takes the complexity out of implementation. It doesn't require months of programmers and consultant's time before you reap the benefits. With AutoPilot, you simply point and click, and within minutes you're in total control.

NASTEL™ TECHNOLOGIES, INC.

Put Your Application Infrastructure On AutoPilot

To avoid such problems, we recommend shifting the focus from immediately allocating functional system requirements to preconceived implementation-level components -- for example, an existing software module or off-the-shelf hardware component -- toward more thorough analysis of an abstraction of the ultimate system.

Requirements analysis, allocation, and validation are clearly critical elements of this approach, but they are not the focus of system analysis and synthesis. The goal is to decompose the system *structurally* in a set of logical steps, following this process:

1. Understand the workings of the system in its intended environment. Treating the system as a black box, identify environmental elements that interact with the system (the system actors) and the actual interactions and results (system use cases and resulting services). In RUP-SE modeling, this is system analysis at the *context level*.
2. Decide how to initially partition the system into elements whose collaborations will provide the services satisfying the needs placed on the system at the context level. Recursively examine elements from both black box and white box perspectives until the key functional and supplementary requirements are satisfied by the interactions of the interdependent system elements. This is investigation at the *analysis level* in RUP-SE.
3. Once the system architecture has been derived abstractly through this process of structural decomposition, the next step is to determine how system elements will be realized. This is a transformation from the abstract to the concrete, with the realization partitioned in some combination of hardware, software, firmware, data, and workers. RUP-SE calls this the *design level* of systems development.
4. Finally, the system is constructed, tested, and prepared for release. The designs specified in the previous step are implemented (optimally in "release-ready" iterative stages as defined within the RUP and RUP-SE frameworks). This final stage of development is captured in RUP-SE's systems development modeling *implementation level*.

As we proceed down from level to level, requirements are more precisely allocated to system elements, uncertainty is progressively addressed and eliminated, and the system can be iteratively demonstrated and assessed. Each level of decomposition forces design decisions, providing ongoing synthesis, coupling of requirements and design specifications, and increasing system detail.

Although the steps above may appear to be sequential, and investigation at a lower, more detailed model level seems to require completing activities at higher levels, this is often not the case in practice, due to a range of factors. These factors are addressed by the next two principles of systems engineering.

Enable both separation and integration of concerns

A natural way to deal with system complexity is to intelligently partition the system into smaller, less complex pieces. The way in which to do this varies, but for similar types of systems, the partitioning tends to follow the same pattern. You might partition an automobile, for example, by grouping related functional components, such as chassis, drivetrain, interior, exterior, electronics/software, and steering/braking/suspension.

However, partitioning a system from only one perspective will not likely facilitate the allocation of all stakeholder needs. Systems engineering is not narrowly focused on building a system satisfying functional needs; rather, it must holistically treat the system in its intended environment as seen by operations, manufacturing, service, system actors, and other stakeholder groups. Performance, ease-of-use, aesthetics, reliability, compliance, and all other key constraints must be addressed. Examining a system from multiple perspectives while simultaneously trying to satisfy dozens of nonfunctional constraints might seem like a way to compound development complexity, not reduce it. Instead of dealing with one perspective with n constraints, the development team has to face m perspectives with n constraints. This is precisely the issue that the second system engineering principle addresses. It is not sufficient to merely partition a system to enable a separation of engineering concerns. The chosen systems engineering development methodology must also provide means for system re-integration. As mentioned earlier, a combination of existing RUP-SE framework elements and the new joint realization table (JRT) construct provide for both the separation and integration of key systems development engineering concerns. This framework is flexible and robust enough to support development activities across a wide range of industries and products.

In practice, the framework is tailored to either a specific product or set of products. The process of tailoring the framework usually occurs early in the development cycle. In RUP, this would be near the end of the *Inception Phase*. The tailoring addresses two primary questions:

1. Which engineering viewpoints must be addressed in system modeling and development? For a joint hardware/software system, this would include at least the *logical*, *physical*, and *information* viewpoints. In most cases, the *geometric* viewpoint would also apply.
2. Which individual model views will be addressed and what artifacts will be developed for each view? For example, the *Analysis Level Logical View* captures the logical decomposition of the system. Here we begin opening the black box that is the system and determining the set of internal elements whose collaboration satisfies the demands actors place on the system, which were established in the *Content Level Logical View*. In our MDSD approach, this logical structure and its accompanying collaborations will be captured using Unified Modeling Language (UML) diagrams in the system model artifact. The product of the MDSD framework tailoring is the development case artifact, which will guide the development team toward tangible results. These results are represented as artifacts in the individual views as development progresses through iterative cycles of system definition, analysis, design, construction, and test.

Specifications flow up and down the architecture

In a perfect systems development

Willing to Buy
Legacy Systems Documentation

User manuals, design documents, software, etc.
describing XML interfaces to CICS 3270
and non-3270 applications.
Documents must have been published before 2002.

Examples of such material might include:
Crossplex from Softouch Systems
Microsoft Babylon & TI-COM
Forte Software
Novera Software
AquiSoft InnerAccess

Send all responses in care of
CICS Documentation
PO Box 2675
New York, NY 10108

Include name and contact information. Responses should also include
identification of the material by title, publication date,
document number, version, etc.

environment, the processes outlined above to decompose a system structurally while enabling both separation and integration of concerns should naturally proceed in a top-down fashion. However, perfection is rare. Uncertainty, misunderstanding, miscommunication, change, technical shortfalls, and other factors all inhibit a top-down flow from context to analysis to design, and ultimately to implementation, from proceeding flawlessly. Acknowledging these inhibitors and addressing them as early as possible in the system development lifecycle is ultimately the key to success and a foundational principle of model-driven systems development.

In reality, specifications flow up and down the architecture. In fact, specifications flow up, down, and sideways. Discoveries made in one view may very well radiate out to affect artifacts within other views on the same model level; they may also flow both up and down model levels within the same viewpoint. Is this an indication that the system's architecture was flawed in the first place? Not necessarily. For instance, discoveries made relatively early in the development cycle, while analyzing system or subsystem prototypes, can easily require significant modifications to the system architecture.

Here's a specific example: Typical printing system architectures include a single subsystem to transform input data into the format required by the marking (imaging) engine. This subsystem normally resides in the printing system's control unit (a single locality). Early in a particular printing system's development, prototype testing revealed that there was insufficient processor, memory, and data bandwidth available to meet specified page printing targets. Based on computing technology available at the time, the only way to meet the bandwidth requirements was to spread the data transformation across three computers and manage the work division with another subsystem. In this case, a technical constraint uncovered with a prototype led to the addition of three localities, associated new connections between localities, new interfaces, and changes/additions to the logical subsystems within the printing system's control unit.

Similar technical considerations apply to many different types of products that require designers to choose where and how to implement function. Thermal constraints could well force design decisions in computers, game consoles, and the like. Geometric constraints (shape and volume) clearly affect design decisions for devices like cell phones; each new product generation has packed more function into ever smaller forms. Physical constraints (mass and density) may also require modifications in machines such as automobiles and airplanes.

Addressing these constraints and others that manifest themselves as system development proceeds is a discovery process. As design decisions are made, it is important to capture their essence, including what effects they will have on the product under development and its model-based representation.

Systems and components collaborate; so should development teams

The first three principles outlined above provide technical guidance for system development teams. They direct us to decompose the system structurally, examining it from multiple viewpoints representing significant areas of concern. They also provide mechanisms to re-integrate the system as a whole, while recognizing that specifications will certainly be discovered not just in the planning and concept stages, but also in design and construction as well. These discoveries will lead to design decisions that may affect more system elements than those contained in the view in which the specification was initially determined, leading to changes up, down, and across the architectural levels.

As systems increase in complexity, greater demands are placed on components and subcomponents. Services are grouped logically to meet both functional and supplementary demands, and the interactions between system entities must be adequately described to allow the system to meet its intended purpose. As the model moves from abstract to concrete elements, functions are allocated to software, hardware, firmware, and workers, and the collaborations between elements are finalized.

In hardware/software codevelopment efforts, collaborations between development teams take on greater importance. Many enterprises align engineers and developers by expertise; software engineers focus primarily on software system elements, hardware engineers may be aligned with functional components (an automobile's engine or chassis, for example), and so on. This alignment can lead to *stovepiping*; that is, narrow bands of specialized knowledge with little interaction across the bands. Although the alignment itself may be a good practice, stovepiping is not. Hence our systems engineering principle: Systems and components collaborate, and so should development teams. This often requires direct management focus, both to encourage and facilitate development team collaborations and to determine the appropriate amount of cross-domain and higher-level resources. An example of this extra management focus might be an enterprise-level systems architecture group encompassing product architectures.

Such an MDSD framework and methodology can aid management in facilitating development team collaboration, as can software tools that automate and enforce the collection and management of system development artifacts. The latter are key to capturing design decisions and provide traceability for requirements derivation and allocation across system elements.

Whenever possible, communication and collaborations should be assessed and realigned as needed to support clear, effective, and timely exchange of knowledge and artifacts. One way to do this (which has been implemented across IBM development) is to form *integrated product teams* composed of key product development stakeholders, and to hold regular team meetings to report status and ensure key stakeholder needs are discussed and addressed throughout the development lifecycle.

Development organizations should reflect product architectures

The previous principle of systems

Motive brings an enlightened approach to application management, building management intelligence into the application itself. Only Motive transcends the fragmented, disjointed reality of traditional application management, to deliver the visibility, insight and automation that support and development teams need to keep applications running smoothly.

This is no mere vision for the future. It's here now in application management products from Motive. Learn more about Motive's breakthrough approach in a new white paper about built-in management at www.motive.com/within1.

motive

www.motive.com

engineering focuses on collaborations between development teams to ensure that interactions between system elements can be fully addressed from pertinent multiple viewpoints.

The next principle is related and clarifies a best practice for optimizing collaborations and communications. It states that development organizations should reflect product architectures.

Development organizations aligned to product architectures are both efficient and effective. The alignment naturally optimizes the communication paths between development teams, and it groups specialized sets of engineers, programmers, testers, and other development personnel together to provide in-depth skills and a deep bench of resources.

Our premise is that successful enterprises with successful products likely have succeeded in this alignment, either implicitly or explicitly. However, over time, markets and products change. When the change involves important shifts in product content and function, the enterprise must react and realign to the changing product architecture. If it doesn't, the required collaborations and communications begin to suffer and lose their effectiveness.

Consider the automotive industry. Until recently, automobiles were largely composed of hardware. Engines and drivetrains, braking systems, chassis, exterior, interior, and safety features like lights, windshield wipers, and emergency brakes were key components and subcomponents. This led to development organizations with bodies of expertise in these areas. When a new product was considered, product teams were formed by picking resources from each of these areas of competency. However, currently software is assuming greater importance in automobiles and contributes heavily to new product content. As more function is implemented in software, increased collaboration and ongoing communication between product teams, hardware domains, and software engineers is needed. This change in content hopefully will be accompanied by refinements in development organization structure to properly reflect software's growing influence.

Base the lifecycle on removing risk and adding value

The above principles collectively address the three key aspects of systems engineering: systems architecture, organizational structure, and process. The three are held together by appropriate governance: program and project management, executive steering, change and configuration management, contractor and business partner management, and so on. The governance team is responsible for ensuring the highest risks are removed first and that product or program value can be measured regularly as development progresses. This leads to the final system engineering principle: Base the lifecycle on removing risk and adding value.

For complex systems, uncertainty is usually the biggest risk a program faces in the early stages of development. A second major risk area is not validating assumptions early in the development process (which may prove to be invalid when the system is in integration testing, leading to significant rework). A third risk area is committing too early to a specific implementation of hardware, software, and firmware, resulting in brittle systems that cannot be easily extended.

How is value measured? Before the methods can be determined, a prerequisite assessment is needed to determine what "value" means on a given program or product. If the measure of value is providing a system that meets stakeholder needs, one approach is to demonstrate the developing system to key stakeholders at regular intervals. This approach is included in both RUP and RUP-SE. Systems are developed iteratively, regularly increasing in capability. The development teams solicit stakeholder review and feedback with each iteration and make appropriate system adjustments in response.

These six principles and the approach they embody are intended to remove risk and provide a framework for complex hardware/software systems codevelopment

YOUR TESTING SUCKS.

1. Writing test code isn't testing – it's coding.
2. Your developers aren't making testable software.
3. It's impossible to get the business side involved in quality.
4. Your existing testing tools have become shelfware.
5. Your customers have become the primary testers.

LISA from iTKO.

Because everyone should own quality.

iTKO Try a free eval of LISA 2.5 at **www.itko.com**

We understand.

You've probably invested considerable time and money trying to automate your testing. But it just didn't work.

You were sold a tool that let you point and click to test on Day 1, and on Day 2 you were maintaining a test script - inserting bugs into your tests at the same rate as you inserted bugs into the code during development.

The tools were too hard for the QA teams to use, they required intense training and developer knowledge, and the seats were more expensive than a presidential fundraiser.

It's OK. We've seen this scenario at every customer we talk to.

LISA is different.

LISA is a comprehensive, no-code automated testing solution ideal for distributed applications and teams.

Using "inline testing" technology, LISA allows even non-developers to instrument and test web services, J2EE apps, databases and websites.

LISA automates test design, staging and monitoring for unit, regression, load and performance tests. And since LISA is pure Java, she's equally stylish on Linux, OS X, Unix or Win clients.

But don't take this ad's word for it. Try LISA for free and talk to the experts at iTKO. We'll show you how everyone can own quality.

Deploy Complete Software Solutions

Leverage IBM Express Runtime

BY LIM SHAM SHAM & TAN ENG KIAN

For independent software vendors (ISVs) and their customers, there is more to a solution than just the application. A complete solution includes an application, but it's an application that is integrated with the middleware components that enable the application to work in the customer's environment. To turn their applications into solutions, it is common for an ISV to have to navigate through the complex process of obtaining, installing, configuring, and supporting many different middleware components from multiple suppliers with different terms and conditions, different licensing structures, and so on. From there, the process of integrating an application with all the different supporting components -- and deploying them all together for the customer -- is often equally complex.

IBM Express Runtime changes all that. IBM Express Runtime offers a cohesive set of middleware components and a single environment that enables ISVs to create and deploy middleware along with their own applications as a single package across multiple platforms quickly and effectively. This article explains what you need to know to begin developing and deploying complete solutions for mid-sized businesses using IBM Express Runtime V2.1.1, including:

- Assembling a Web-based Enterprise JavaBeans (EJB) application.
- Using three of IBM's main middleware components -- IBM WebSphere Application Server - Express V6.0.2, IBM DB2 Universal Database Express Edition V8.2 Fixpack 3, and IBM HTTP Server V6.0.2 -- all of which are included with IBM Express Runtime V2.1.1.
- Using the Express Runtime Developer.
- Deploying the integrated solution using the deployment wizard.

We will look at IBM Express Runtime

(This article was first published on developerWorks WebSphere at http://www.ibm.com/developerWorks/websphere.)

ABOUT THE AUTHORS

Lim Sham Sham works as a technical support representative with Virtual Innovation Center supporting DB2 and Express Runtime. She is an IBM Certified Database Administrator and Application Developer for DB2 UDB V8. She is also an IBM Certified Application Developer for WebSphere Studio Application Developer.

Tan Eng Kian works as a technical support representative with the ISVs and Developer Relations unit supporting WebSphere and Rational Software. He is an IBM Certified System Administrator and Application Developer for WebSphere Application Server and WebSphere Studio Application Developer.

and walk through the development and deployment process in these sections:

1. IBM Express Runtime product and process overview
2. Develop a Web-based EJB application
3. Create application wrapper and package with the required middleware components
4. Test the solution wrapperThis article assumes a working knowledge of J2EE application development. To follow the steps in this article, you will need to have IBM Express Runtime V2.1.1 properly installed in your environment.

Visit the IBM PartnerWorld Web site for details on obtaining an evaluation copy of IBM Express Runtime..

IBM Express Runtime product and process overviewLet's take a quick look at what IBM Express Runtime is and what you can do with it.

IBM Express Runtime V2.1.1 (hereafter called Express Runtime) is made up of the following main components:

- Fully functional IBM middleware components that you need to run applications for mid-sized businesses:
  - IBM WebSphere Application Server - Express V6.0.2
  - IBM DB2 Universal Database (UDB) Express Edition V8.2, Fixpack 3
  - IBM Informix® Dynamic Server Express V10.0 TC3E/UC3E
  - IBM HTTP Server V6.0.2 o Web server plug-in for WebSphere Application Server - Express 6.0.2
- Solution assembly toolkit consisting of:
  - Express Runtime Developer that helps you assemble your application, along with IBM middleware components.
  - Deployment wizard to quickly and effectively deploy your solution to a target system.
- IBM Rational® Web Developer V6.0 application development tool
- Express Runtime console provides a single, Web-based utility for managing IBM middleware components.

In general, there are five steps involved in developing-to-deploying a solution in the Express Runtime environment, as shown in Figure 1.

Step 1: Use IBM Rational Web Developer V6.0, or equivalent application development tools, to develop your own J2EE applications.

Step 2: Install Express Runtime on an internal development machine. The following artifacts are included with the Express Runtime installation:

- Solution Assembly Toolkit, which consists of Express Runtime Developer and Deployment Wizard
- IBM middleware images
- IBM Rational Web Developer V6.0
- Express Runtime console
- Documentations and samples.

Step 3 consists of two activities:

- Integrate your applications with the required IBM middleware components using the Express Runtime Developer tool. A solution that consists of your application and all the required middleware components will be created.
- The solution is then stored (on CD, DVD, or on a staging server) for later deployment to customer environment.

Step 4: Use the Deployment Wizard to deploy the stored solution.

Step 5: Manage IBM middleware components through the Express Runtime console, a Web-based console that provides consolidated administration of WebSphere Application Server - Express and DB2 UDB Express, or Informix Dynamic Server and HTTP Server.

Develop a Web-based EJB applicationThe goal is to create not just a valuable application, but a complete packaged solution that, upon installation, is already integrated with the components it needs.

We will walk through the process, beginning with a completed EJB application. For the purpose of this article, we will use a Personal Budget Analysis (PBA) EJB application as an example.

To create a packaged solution with this application:

- We will use three middleware components to run the application: WebSphere Application Server - Express, DB2 UDB Express, and HTTP Server.
- We will use Express Runtime Developer to package the PBA application together

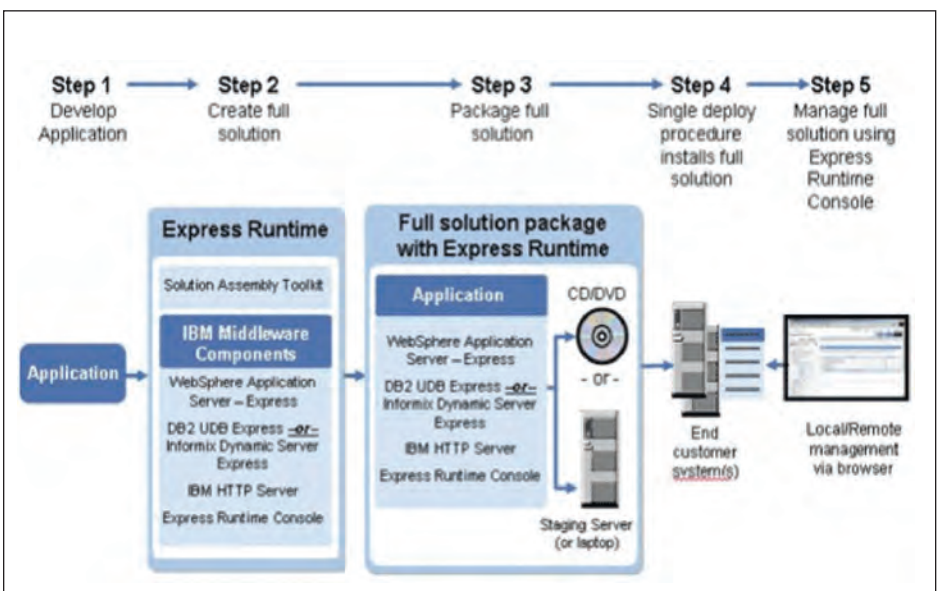

Figure 1. Express Runtime architecture

with these middleware components into a solution.
- We will deploy the solution to a customer environment using the Deployment Wizard.

About the sample applicationHere is some background on the sample application we will be working with.

From a business function standpoint, the PBA application captures income and expense data and generates an income-and-expense ratio graph. The application was developed using IBM Rational Application Developer V6.0 with the UML editor as the development tool, along with JSPs, servlets, container-managed persistence (CMP) entity beans, session facade, and Service Data Objects. Figure 2 shows you the relationship diagram for CMP entity beans used in our application.

Every individual user obtains income and incurs expenses:
- Income consists of salary income and non-salary income.
- Expense consists of household expenses, repayment for loan borrowing, and payment for insurance premium.

In this model, we use an entity bean named PBUser which represents a user in the Personal Budget database. PBUser has one-to-one relationship reference with the following entity beans:
- Income
- HouseholdExpense
- BorrowingAndLoan
- InsurancePayment.

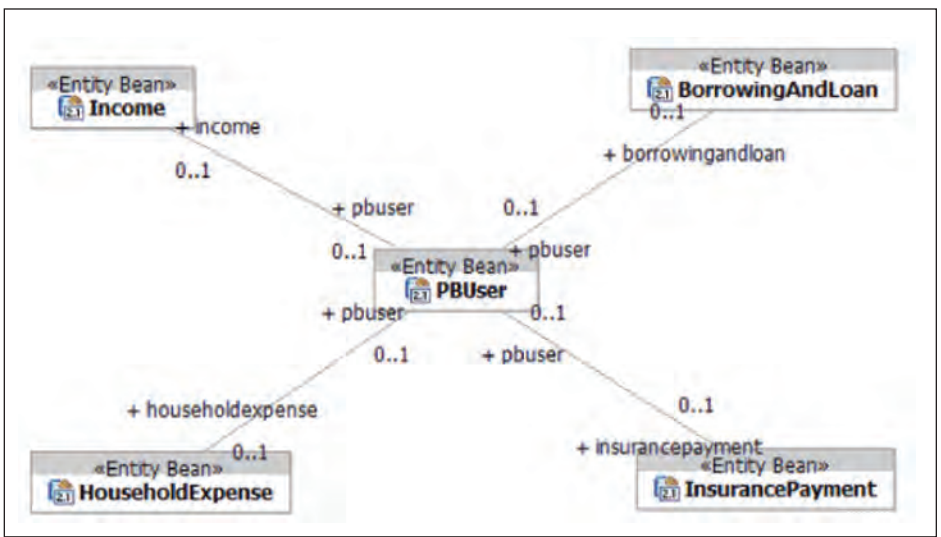

Figure 2. Relationship diagram for CMP entity beans

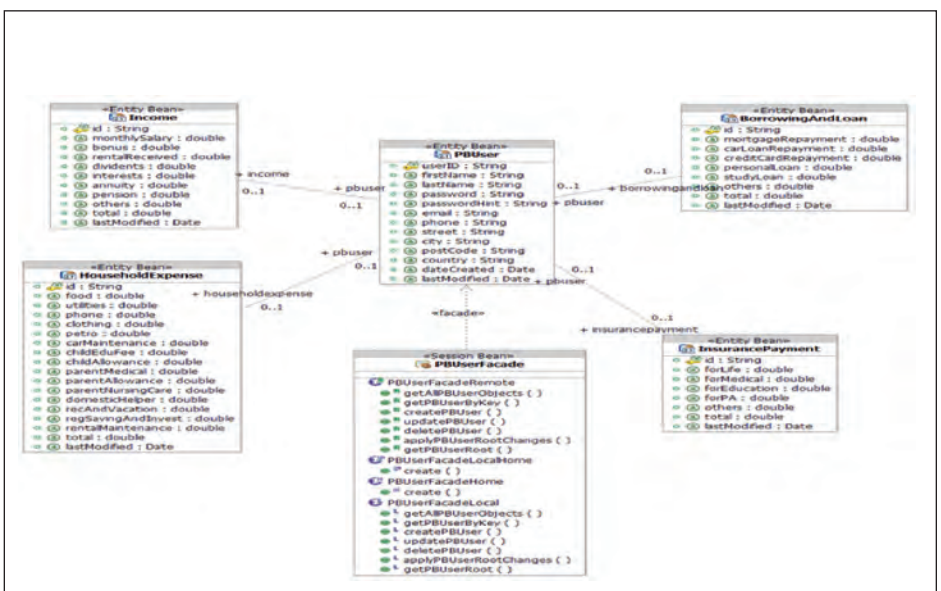

Figure 3. Class diagram for PBA application

Figure 3 shows a class diagram depicting the entities and their relationships. (This class diagram was created using the UML visual editor in Rational Application Developer.)

The class diagram shows the attributes for each entity bean. You will also notice a session facade bean named PBUserFacade, which contains the create, update and delete (CRUD) APIs to manage the instances of the PBUser SDO, which was added to the facade using the value-objects attribute. The methods it creates (shown in Figure 3) are getAllPBUserObjects, getPBUserByKey, createPBUser, updatePBUser, and deletePBUser. This session bean will be used by an ActionServlet in PBudgetWeb module to interact with PBUser entity bean to manage the user instance and its related entity bean instances.

The typical process flow for the PBA application is:
1. New user signs up as a user by completing the user registration page.
2. Registered user logs on to the PBA system using their user ID and password.
3. Registered user performs one of the available personal budget functions:
   - View the Income and Expenditure Budgeting summary page (Figure 4).
   - View or enter details of personal monthly income (Figure 5).
   - View or enter details of monthly household expenses (Figure 6).
   - View or enter details of monthly loan repayment expenses (Figure 7).
   - View or enter details of monthly insurance premium payment (Figure 8).

Create PBA application wrapper and package with the required middleware componentsBefore assembling the application, we need to understand some key

Visit the *New* www.SYS-CON.com Website Today!

The World's Leading *i*-Technology News and Information Source

24/7

JUMP TO THE LEADING i-TECHNOLOGY WEBSITES:

IT Solutions Guide
Information Storage+Security Journal
JDJ
Web Services Journal
.NET Developer's Journal
LinuxWorld Magazine
Linux Business News
Eclipse Developer's Journal
MX Developer's Journal
ColdFusion Developer's Journal
XML Journal
Wireless Business & Technology
Symbian Developer's Journal
WebSphere Journal
WLDJ
PowerBuilder Developer's Journal

FREE NEWSLETTERS
Stay ahead of the i-Technology curve with
E-mail updates on what's happening in your industry

SYS-CON.TV
Watch video of breaking news, interviews with industry leaders, and how-to tutorials

BLOG-N-PLAY!
Read web logs from the movers and shakers or create your own blog to be read by millions

WEBCAST
Streaming video on today's i-Technology news, events, and webinars

EDUCATION
The world's leading online i-Technology university

RESEARCH
i-Technology data "and" analysis for business decision-makers

MAGAZINES
View the current issue and past archives of your favorite i-Technology journal

INTERNATIONAL SITES
Get all the news and information happening in other countries worldwide

terminology as it is used within the Express Runtime context:

- Wrappers can be thought of as the "workflow" or "controller" elements that are responsible for controlling the flow of execution for deploying a component. There are two types of wrappers available:
  - Application Wrapper
  - Solution Wrapper.
- An application is a component of an overall solution; it can be a middleware component or an actual application. For example, our PBA application is considered an application wrapper in the context of Express Runtime; DB2 UDB Express -- which is one of the components required to support the PBA application -- is also an application wrapper.

There are some ready-to-use sample application wrappers provided in the Express Runtime Developer tool that you can use as templates for your own application wrappers. Each application wrapper contains information specific to the application, such as name and version of the application, user programs to be used by the application, and so on. Application wrapper information is kept in a XML file called application.axml, which can be found in the src folder in the Express Runtime Developer tool.

- When your application wrapper development is complete, you can then package all the required application wrappers into a single solution wrapper; a solution wrapper consists of one or more application wrappers.

The Express Runtime Developer tool provides some ready-to-use sample solution wrappers, each of which contains information specific to the solution, including the ID, version, supported language of the solution, and so on. Solution wrapper information is kept in a XML file called solution.sxml, located in the src folder in the Express Runtime Developer tool, as shown in Figure 10.

The solution we will be developing, PBASolution, is made up of five application wrappers:

- PBA
- WebSphere Application Server - Express
- DB2 UDB Express
- HTTP Server
- Web server plug-in.

Since WebSphere Application Server - Express, DB2 UDB Express, HTTP Server, and the Web server plug-in are part of the sample application wrappers provided in the Express Runtime tool, we only need to develop the application wrapper for PBA. Subsequently, all of these application wrappers will together be packaged into a solution wrapper, as shown in Figure 11.

With an understanding of wrappers and how they are structured, we will move on to the next level of PBA application wrapper development. More terminology will be introduced along the way.

To develop our PBA application wrapper, we will perform these general tasks, each of which is explained in detail in the following sections:

A. Create PBA application project
B. Edit Application Wrapper
C. Create Budget.prop response file
D. Create PBA user programs
E. Add wsadmin and DB2 scripts
F .Complete PBA application project
G. Create PBASolution solution project
H. Edit Solution Wrapper Editor
I. Complete PBASolution solution project

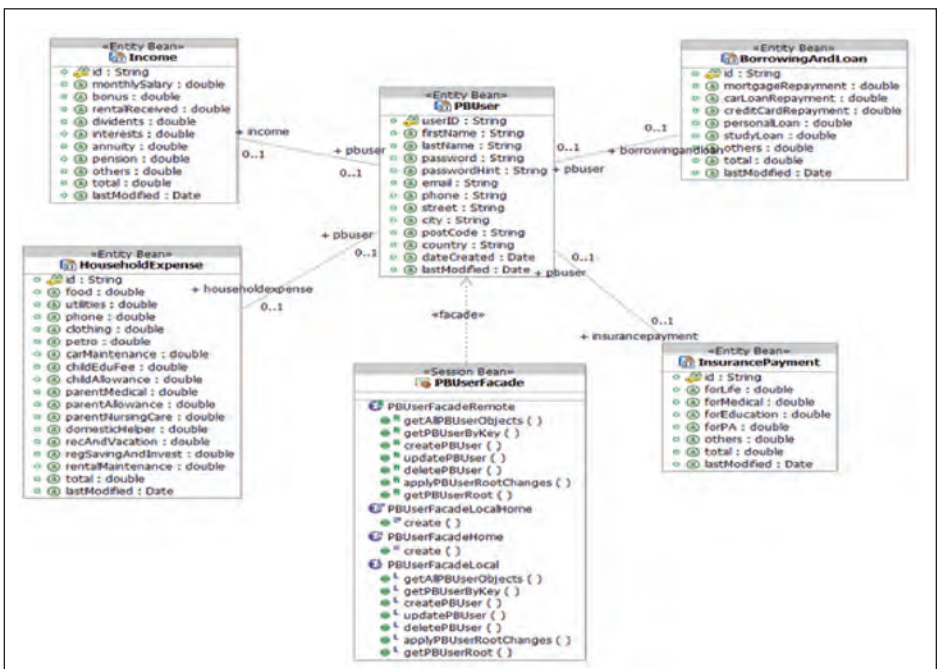

Figure 4. Summary page for Personal Income and Expenditure Budgeting

A. Create PBA application project

1. Start Express Runtime Developer by selecting Start => Programs => IBM Express Runtime 2.1 => Express Runtime Developer. You should now be in the Express Runtime Developer perspective.
2. Create a new application project by selecting File => New => Express Runtime Application Project.
3. On the next dialog, enter PBA for the Project name and select Next.
4. Next, enter or select the following values use the following values:
   - Application ID: PBAo Version: 1.0
   - Installation time (minutes): 30o Operating system: Windows
   - Default language: English
   - Wrapper file name: application.axmlSelect Finish to complete the creation of PBA application project.
5. From the Package Explorer panel, you will see that the PBA application project has been created and that the Application Wrapper Editor Welcome page is displayed in the right pane (Figure 12).

Reach Over 100,000

Enterprise Development Managers & Decision Makers with...

Offering leading software, services, and hardware vendors an opportunity to ***speak to over 100,000 purchasing decision makers*** *about their products, the enterprise IT marketplace, and emerging trends critical to developers, programmers, and IT management*

Don't Miss Your Opportunity to Be a Part of the Next Issue!

Get Listed as a Top 20* Solutions Provider

For Advertising Details Call 201 802-3021 Today!

*ONLY 20 ADVERTISERS WILL BE DISPLAYED. FIRST COME FIRST SERVE

SYS-CON MEDIA

The World's Leading i-Technology Publisher

B. Edit Application Wrapper

Next, we will use the Application Wrapper Editor to edit the PBA application wrapper.

The **General** tab shows the details you specified above, which includes the general information of the application project. You can use this page to select the appropriate options; since our sample application only runs on Windows and the only supported language we selected is English, these are the only options you will see selected.

The **Programs** tab gives you the information about any custom user programs, which can be used to:

- Perform pre-installation, installation, or post-installation of IBM middleware components and ISV applications.
- Execute wsadmin scripts written in Jython or JACL to configure the WebSphere Application Server environment.
- Execute DB2 scripts to create databases or tables, and populate tables with data.

(Rational Web Developer, which is included with Express Runtime, can be used to develop Java-based user programs.)

There are four types of user programs that can be developed for use with Express Runtime, differentiated by the point during the deployment process in which the program runs:

- A predeployment checker is run prior to the entry, main, or exit programs, and is used to examine whether target computer meets certain criteria, such as having adequate disk space, any conflicting software, a previous version of the application, and so on.
- An entry program is typically used to perform any pre-configuration or setup required before the main installation program runs.
- The main program is the installation program that performs the installation of your application(s) on target computer.
- An exit program is typically used for any post-configuration required following the product installation, and can be used to determine if the installation of the solution is successful.

Our sample application project will use two user programs: a predeployment checker and a main program. The predeployment checker will be used to determine whether the target computer has all necessary prerequisites in place, and the main program will perform the actual installation of our PBA application, which will include executing DB2 scripts, creating a DB2 database and tables, installing the PBA application on WebSphere Application Server - Express, and invoking wsadmin scripts to configure WebSphere Application Server - Express environment resources.

You may notice at this point that there is an error in the PBA application project. This is because a main program is required for every application project, but, by default, no program is initially associated with it. This error will disappear after the main program has been defined.

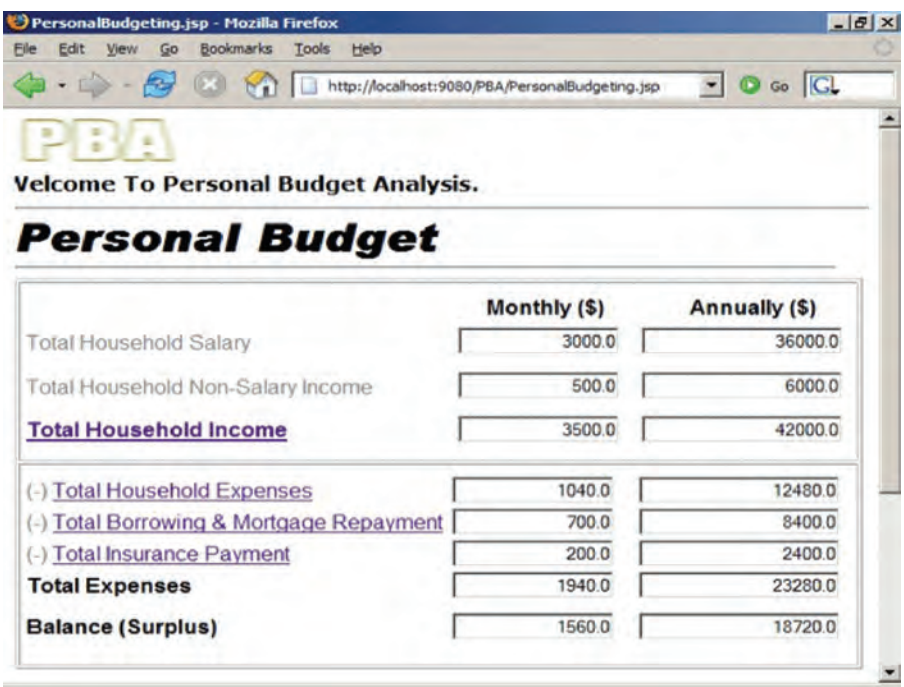

Figure 5. Details for household income

To define both of the user programs we need:

1. Select the Programs tab to be taken to the Main Program page.
2. In the Basic Program Configuration section, make sure that the Program type field is set to Java program.
3. In the Program field, type in com.budget.PBAWinMain. This will associate the PBAWinMain class as the Main Program.
4. In the Program Arguments section, enter Budget.prop into the Response file field, and PBAWinMain.log into the Log file name field.
5. In the Arguments field, select Add, from the Add Argument window, select Response file name, and then click Next.
6. Make sure the Response file name field is set to Budget.prop and then Finish. The PBAWinMain program will take Budget.prop as an input argument.

We have completed the definition of the Main Program, as shown in Figure 13. Next, we will define the predeployment checker user program.

7. Change to the Predeployment Checker page and check Enable predeployment.
8. Make sure Java program is selected as the Program type in the Set Program

window and click Next.

9. In the Specify the main class window, enter com.budget.PBAWinPDC in the Main class field and click Finish.
10. In the Program Arguments section of the User Programs Information window, enter Budget.prop in the Response file field and PBAWinPDC.log in the Log file name field.
11. In the Arguments field, click Add.
12. From the Add Argument window, select Response file name and click Next.
13. Make sure the Response file name field is set to Budget.prop and click Finish.
14. Save the changes we have made so far by selecting File => Save from the main menu.

We have completed the definition of our predeployment checker, our second user program, as shown in Figure 14.

The **Variables** tab enables you to configure the variables you need in deploying applications. For our sample application project, two variables, DB2UserId and DB2UserPassword, are required. We will first add the variables and then set the following validation rules for these two variables (you can always add more, if necessary):

- Valid characters: @#$_abcdefghijklmnopqrstuvwxyz0123456789
- Invalid values: ADMINS, USERS
- Invalid prefixes: SQL, _, IBMTo configure these variables:

15. In the Application Variables section, select Add.
16. Select String Variable as the Variable type. Enter DB2UserId into the Variable name field and DB2 Administrator User ID into the Variable description field. Click Finish.
17. You should now see a string variable called DB2UserId added in the Application Variables list.
18. In the Variable Validation Configuration section at the bottom of the window, define the following:
    - Default value: db2admino Minimum length: 2
    - Maximum length: 3019.

Click Add next to the Validation rules field to add each of the validation rules listed above. To define the first validation rule, select Specify a valid value in the Add a variable validation rule window, then Next.

20. Make sure you select the Specific characters option and then Next.
21. In the next window, enter the following in the Valid string field: @#$_abcdefghijklmnopqrstuvwxyz0123456789 and then Finish.
22. Select Add next to the Validation rules field again to define the second validation rule.
23. In the Add a variable validation rule window, select Specify an invalid value and click Next. In the next window, keep the Complete String default option and then Next.
24. Enter the following in the Invalid string field: ADMINS and click Finish.
25. Repeat steps 22 through 24 to add an invalid string for USERS.
26. To define an invalid prefix, click Add next to the Validation rules field again.
27. In the Add a variable validation rule window, select Specify an invalid value, then Next.
28. In the next window, make sure you select the Prefix option and then Next.
29. Enter SQL in the Invalid prefix field and click Finish.
30. Repeat steps 27 through 29 to add an invalid prefix for "_" and "IBM".

You may notice a section called Variable Associations Configuration. This option is used to perform variable associations, making variable values available to user programs by including the values in the response file (properties file). We will perform variable associations for DB2UserId in the next set of steps, making the value of DB2UserId available to user programs through the Budget.prop properties file.

There are few different types of associations available. We will be using Properties association in our example. All Java-based user programs can access the variable values through this type of association.

31. Click Add to the Variable Associations Configuration section.
32. In the Add Association window, select Properties association and then Next.
33. Enter DB2UserId in the Keyword field and click Finish.
34. You should now have DB2UserId in the Variable Associations Configuration list.
35. Save the changes by selecting File => Save from the main menu.

We have completed defining our first variable, DB2UserId, as shown in Figure 15. Next, we will define our other variable.

The second variable we will add is DB2UserPassword, which has the following validation rules:

- Valid characters: @#$_abcdefghijklmnopqrstuvwxyz0123456789
- Invalid values: ADMINS, USERS
- Invalid prefixes: SQL, _, IBMTo add the DB2UserPassword variable:

36. In the Application Variables section, click Add.
37. Select Password Variable as the Variable type. Enter DB2UserPassword for the Variable name, and DB2 Administrator User Password as the Variable description, then Finish.
38. DB2UserPassword is now added in the Application Variables list.
39. In the Variable Validation Configuration section, define the following values:
    - Minimum length: 2
    - Maximum length: 3040.

Follow the same steps for DB2UserId to add the validation rules for DB2UserPassword.

41. Click Add next to define the Variable Associations Configuration for DB2UserPassword.
42. In the Add Association window, select Properties association, then Next.
43. Enter DB2UserPassword in the Keyword field and click Finish.
44. DB2UserPassword should now appear in the Variable Associations Configuration list.
45. Save the changes by selecting File => Save from the main menu.

The completed variable definition for DB2UserPassword as shown in Figure 16.

The **Files** tab enables you to define all the files needed for installing your application. We will come update this tab later, after we have created all the required files (next).

(To be continued next issue.)

BONUS TRACK ADDED! (FREE $1,295 VALUE!)
"Real-World AJAX" 2-Day Seminar

SOA WebServices & Enterprise OpenSource Conferences

The #1 i-Technology Educational and

Register at... www.SOAEOSConference.sys-con.com

The **10th International SOA Web Services Edge Conference**, colocated with the First Annual **Enterprise Open Source Conference**, will take place on **June 5-6, 2006** at the historical **Roosevelt Hotel in New York City**. SOA Web Services Edge and Enterprise Open Source Conference together deliver the #1 i-technology educational and networking opportunity of the year.

This year, more than 100 distinguished conference faculty members will present **96 cutting-edge SOA, Web Services and Enterprise Open Source topics** in educational classes, panels, and keynotes in six simultaneous tracks for two information-packed days, plus a bonus track, a **"Real-World AJAX"** two-day seminar, a $1,295 value!

The conference program does not merely present a comprehensive view of all the development and management aspects of integrating a **SOA strategy** and an **Open Source philosophy** into your enterprise, its organizing principle is that delegates will go away from the intense two-day program replete with why-to and how-to knowledge delivered first-hand by industry experts.

According to analyst firm Gartner Group, **by 2008 more than 60 percent of enterprises will use SOA** as the guiding principle when creating mission-critical applications and processes. **"Businesses that ignore the potential of SOA will find themselves outpaced by rivals who improve their agility and transform themselves into new kinds of enterprises,"** says Gartner analyst Yafim Natis.

Meanwhile, the **Open Source revolution** has spread into traditional areas of enterprise IT, with all major vendors and a global group of **Open Source entrepreneurs** developing the technology and delivering the solutions that are changing the way the world does business and the way governments serve their citizens.

Receive FREE WebCast Archives of all Entire Conference!

The best news for this year's conference delegates is that your "Golden Pass" registration now gives you full access to all 96 conference sessions. We will mail you the complete content from all the conference sessions in seven convenient DVDs after the live event takes place.

This on-demand archives set is sold separately for $995

Receive 8 Days Worth of Education F

With WebSphere Process Server

Integrate EJB Services

BY PETER XU

IBM WebSphere Process Server is a runtime platform for business integration solutions developed using IBM WebSphere Integration Developer.

Many existing business functions for a company's IT infrastructure are written based on Java™ 2 Enterprise Edition (J2EE) stateless Enterprise Java Beans (EJB). When you design and implement new business process integration applications, it is important that you have the ability to leverage and integrate these functions easily.Overview diagramThe scenario we use in this article is pretty straightforward. It consists of one Service Component Architecture (SCA) module that has a Java component, which uses the imported EJB service.
Software prerequisitesYou need to install:

- WebSphere Integration Developer V6
- WebSphere Process Server V6 test environment

This article assumes that you have some knowledge of J2EE and have used WebSphere Studio Application Developer or Rational Application Developer before. To learn more about these products, visit developerWorks.

Configuration

In this section, we develop a sample EJB application and an SCA module that utilizes the EJB service.

You can begin by creating the enterprise application, package, and EJB session.

1. Create an enterprise application called Echo with one EJB module, EchoEJB2. Create a package called com.ibm.issw.poc.invokeejb.
3. Create a stateless EJB session called Bean Echo
4. Create a simple method called echoSimple, and promote it the remote interface:

```
public String echoSimple(String name){
    System.out.println("inside echoSimple");
    return "Hello, " + name;
  }
```

5. Open the ejb deployment descriptor, and notice the default JNDI name for the EJB:

```
ejb/com/ibm/issw/poc/invokeejb/EchoHome
```

6. Create an EJB client project and jar that will be used later in the SCA module.
7. Right-click and select EchoEJB => EJB Client Jar => Create EJB Client Project and enter EchoEJBClient in the name field.
8. Run and test the EJB application (see Figure 5).

Now, that we've created the EJB, let's integrate it.

Create an EJB component

To integrate the EJB into the WebSphere Process Server SCA component:

1. Create a module called InvokeEcho.
2. From the Dependencies editor, select EchoEJBClient as the dependent Java project and click Add to add EchoEJBCLient as the dependent Java project.

It is important to add module dependency at this point. When we add the interface later, the wizard will be able to find the remote EJB interface.

3. From the Assembly editor, open the module, InvokeEcho in Assembly editor, add an Import component on the canvas. Change the name to EchoEJB.
4. Add the interface for the import:
   1. Click the Add Interface icon.
   2. In the resulting window, select Show Java.

*(This article was first published on developerWorks WebSphere at http://www.ibm.com/developerWorks/websphere.)*This article shows you how to use IBM WebSphere Process Server Version 6 and IBM WebSphere Integration Developer Version 6 to integrate existing J2EE applications without making changes to them.

ABOUT THE AUTHOR

Peter Xu is a Senior Consultant with IBM Software Services for WebSphere group. Consultants with Software Services for WebSphere help customers deploy IBM products into their organizations. Peter provides consulting services, education, and mentoring on J2EE technologies, and specifically WebSphere and WebSphere Studio products to Fortune 500 clients. You can reach him at peteryxu@us.ibm.com.

Figure 1: Create an EJB

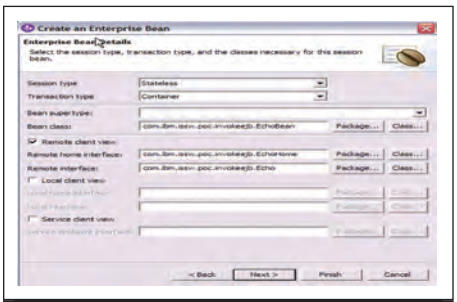

Figure 2: Create an EJB

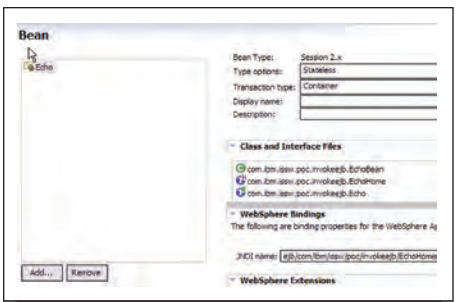

Figure 3. EJB JNDI name

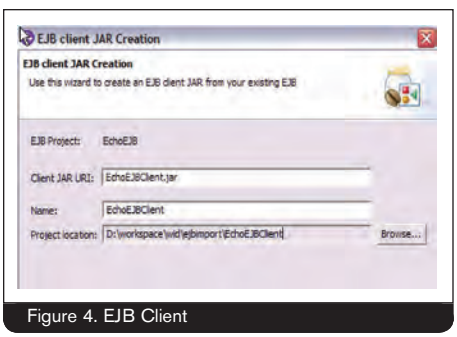

Figure 4. EJB Client

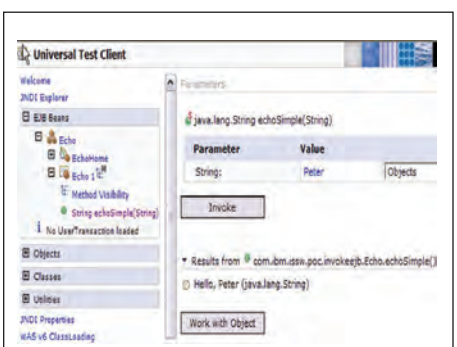

Figure 5. Testing the EJB

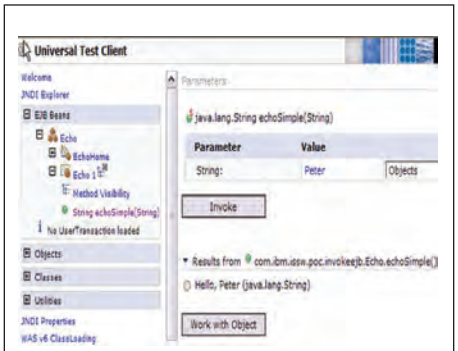

Figure 6. Add dependencies

Figure 7. Add the Import

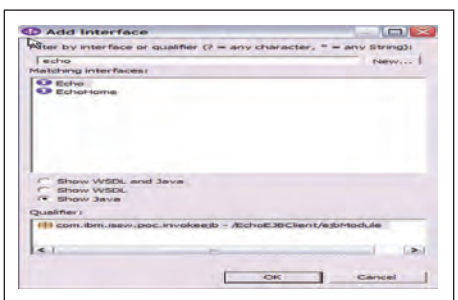

Figure 8. Add Interface for Import

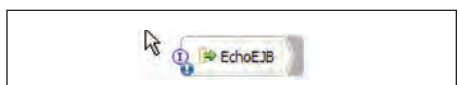

Figure 9. Import with the interface

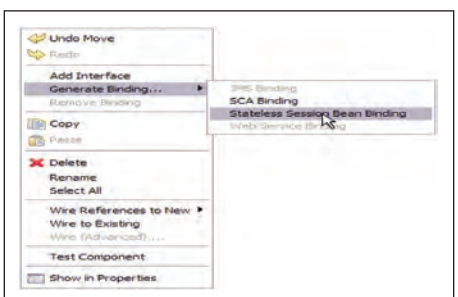

Figure 10. Add Binding

Figure 11. Import with Interface Binding

Figure 12. Specify JNDI name for Binding

Figure 13. Add the Java Component

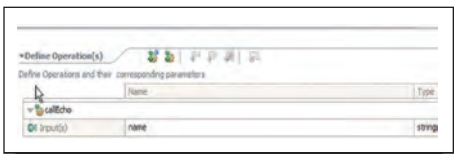

Figure 14. Create an Interface

Figure 15. Wire Java and Import the Component

3. Enter echo in the Filter by interface field.
4. Select Echo from Matching interfaces . This interface is from project /EchoEJB-Client/ejbModule (see Figure 8).
5. Click Ok and save the module diagram.
6. Generate the Session EJB Binding.

1. Right-click the EchoEJB import and select Generate Binding => Stateless Session Bean Binding.
2. Notice how the icon changed in the import (see Figure 11).
3. Select the Properties view and the Binding tab. You can see that the JNDI name field is now populated (see Figure 12).

Create a Java Component

To create a Java Component:

1. Drag and drop a Java Component onto the assembly canvas and change the name to InvokeEcho (see Figure 13).
2. Create an interface called InvokeEchoInterface with a one-way operation callEcho and one input parameter:
3. Select the Java component and add an interface for it, InvokeEchoInterface.
4. Wire the InvokeEcho and EchoEJB together (see Figure 15).

Notice that a reference EchoPartner has been generated. We will use it in the Java

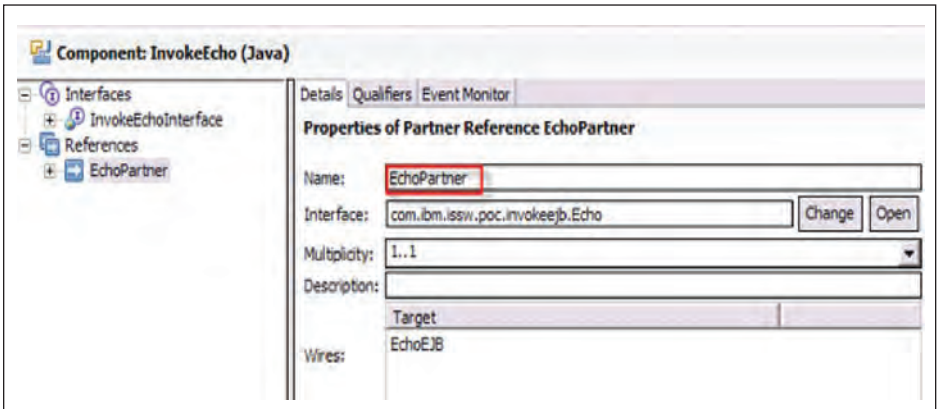

Figure 16. Partner reference

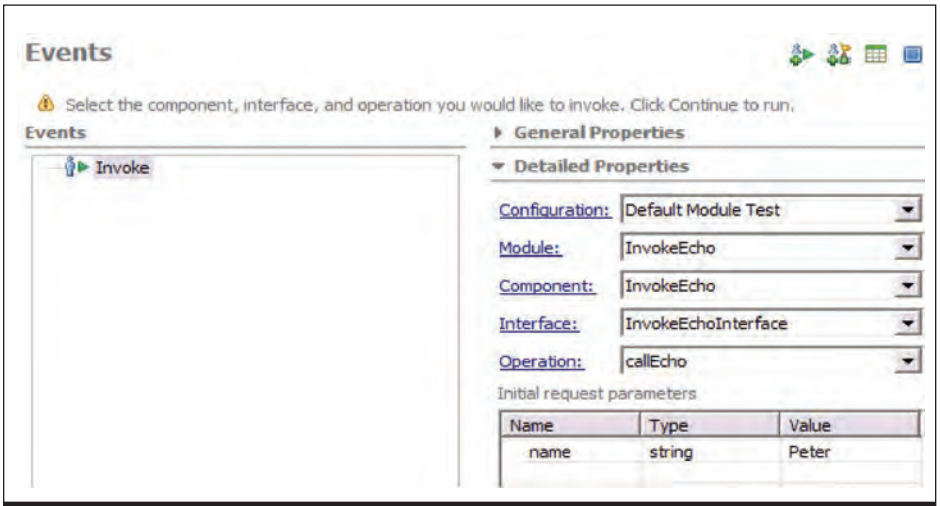

Figure 17. Component Test 1

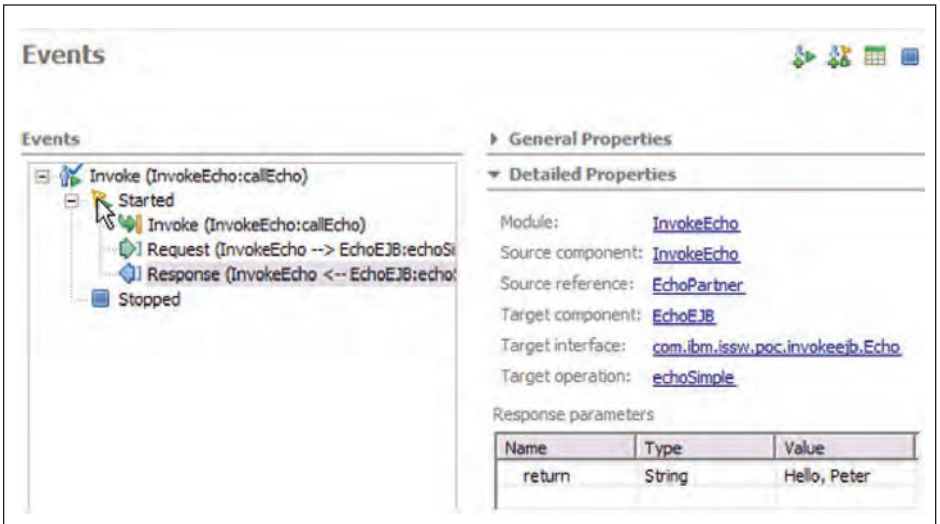

Figure 18. Component Test 2

Component implementation (see Figure 16).

5. Select and right-click the Java Component and generate the default Java implementation. When the implementation opens, replace the callEcho method with the following code:

```
        try {

            System.out.println("in callEcho");
         // Create an instance of the
ServiceManager
           ServiceManager serviceManager = new
ServiceManager();

           Echo echo= (Echo)serviceManager.locateS
ervice("EchoPartner");

           String result = echo.echoSimple(name);
           System.out.println("called imported
EJB, and response is " + result);

        } catch (Exception e) {
            e.printStackTrace();
        }
```

6. Running and testing the scenarioTo test the scenario:

1. Right-click InvokeEcho and select Test Component.
2. In the test configuration, remove the emulator for EchoEJB, since we want to invoke the real EJB implementation.
3. Enter a value for the name field and click Continue.
4. Notice the expected result.

Notice the result from the System.out.

Congratulations, you've successfully imported and invoked the Echo EJB from the SCA component, InvokeEcho. In this scenario, the SCA module and the EJB service are running on the same server.

Conclusion

This article helped you implement a WebSphere Process Server solution that integrates EJB services. As you can see, using EJB services from an SCA component is easy to do. Reusing EJB-based services in this way creates tremendous opportunities for integration products, such as WebSphere Process Server..

WebSphere JOURNAL | Advertiser Index...

COMPANY	URL	PHONE	PG
AJAX Seminar	http://www.ajaxseminar.com/	201-802-3022	48-49
Eclipse	http://eclipse.sys-con.com	201-802-3000	47
Instantiations	www.instantiations.com/analytix	800-808-3737	2
ISSJ	http://issj.sys-con.com/	888-303-5282	41
IT Solutions guide	http://itsolutions.sys-con.com/	201-802-3021	21
iTVCON	www.iTVcon.com	201-802-3021	35
ITKO	www.itko.com	877-BUY-ITKO	15
JDJ	http://jdj.sys-con.com/	201-802-3021	33
Motive	www.motive.com/within1	512-531-2527	13
Nastel	www.nastel.com/uncomplicate	800-580-2344	9
OPNET	www.opnet.com	240-497-3000	7
Parasoft	www.parasoft.com/jdjmagazine	888-305-0041	51
Prolifics	www.prolifics.com	800-675-5419	3
Proskauer			11
Quest	www.quest.com/911	240-497-3000	52
SOA Conference	www.soaeosconference.sys-con.com	201-802-3000	24-25
sys-con media website	www.sys-con.com	888-303-5282	19
Tangosol	www.tangosol.com	617-623-578	2
webservices edge/open source	www2.sys-con.com/events	201-802-3066	37
websphere journal	www.sys-con.com	888-303-5282	
Wily	www.wilytech.com	888-GET-WILY	5

Many Layers Mean Rich Functionality

Under the Hood of IBM Workplace Collaboration Services

BY BOB BALABAN

Take a peek under the hood of IBM Workplace Collaboration Services and learn about the layers that make up Workplace Collaboration Services, including IBM WebSphere Application Server and IBM WebSphere Portal.

ABOUT THE AUTHOR

Bob Balaban is an Executive Consultant with the Business Transformation Team in the Workplace, Portal and Collaboration division of IBM Software Group. Previously, Bob worked on NOI, LotusScript, and agents in Lotus Notes initially as a Lotus employee, and then later as a member of Iris Associates. Before Lotus Notes, he worked at Lotus on spreadsheets and other products, earning a U.S. patent for the Version Manager feature in 123/W.

Long-time users of IBM Lotus Notes and Domino (especially those using them for more than just email) know full well the richness of functionality and the power of the platform for creating all kinds of collaborative applications. Many Notes/Domino fans are not, however, sure what benefits they can derive from combining the power of their Domino-based applications with IBM's J2EE (Java 2, Enterprise Edition) frameworks and product offerings. Some people are still not sure what the real differences are among the three major J2EE products: IBM WebSphere Application Server, IBM WebSphere Portal, and IBM Workplace Collaboration Services. (There are many IBM product offerings under the Workplace umbrella; this article focuses on Workplace Collaboration Services and on the architecture on which it is built).

The purpose of this article is to explain how these three J2EE products are actually layered, one upon the other (Workplace Collaboration Services on WebSphere Portal on WebSphere Application Server) to form a coherent whole. We also present an overview of how Lotus Domino might integrate with each of these layers and what the different problems are that you might solve with each type of integration.

Furthermore, although most people think about server-side integration when talking about Lotus Domino and J2EE, we hope to sketch out how the Notes client will evolve over the next year or two to become a J2EE rich client as well.

Technical Foundation: J2EE and WebSphere Application Server

Let's start with a question many people are embarrassed to ask: What is J2EE? It is a specification owned by Sun Microsystems for a Java-based platform or framework on which people run Web-oriented applications. J2EE is not a product (although many vendors, such as IBM, have implemented products that conform to the spec). It's definitely not a religion -- it's a tool, which is sometimes appropriate to use.

The whole idea of J2EE is to provide a vendor-neutral (and therefore, system-independent) way of hosting Web applications. It gives you an application meta-model (meaning, a set of components and tools that you use to build applications) and a platform consisting of a bunch of services that help you run your application, expressed as a set of Java APIs.

The platform is architected to provide performance and scalability to the applica-

tions that you host on it. Performance is fairly obvious: You want the applications that run on a J2EE foundation to perform well. Scalability is a little bit trickier, but it generally means that you have ways of adding load to the system (for example, add more users or increase the traffic flow) without seriously degrading performance or response time. How does J2EE in general, and WebSphere Application Server in particular, achieve this? Take a look at Figure 1, which provides a high-level architectural view of a basic J2EE application server.

The diagram looks a bit overwhelming at first, but the major components are circled in different colors to show you how the big picture can be broken down into more easily explainable pieces.

The green circle on the left-hand side, enclosing the boxes labeled Applet container and Application client container, represents the client side of the system. These two boxes connect to the Web Container box via HTTP and SSL arrows, indicating that Java clients (either stand-alone Java applications or browser-based applets) use the HTTP wire protocol to communicate with J2EE platforms. It's the platform's job to translate specific application URLs into application component "addresses" and to route the commands properly.

The actual business components that developers write are shown circled in orange. JavaServer Pages (JSPs) and servlets are directly addressable with URLs. They typically process inputs of various kinds, perhaps querying some backend system for data, and then renders a result back to the requesting client, often in HTML or XML formats. JSPs and servlets run in the context of a Web Container, and a collection of these--together with image resources, other Java classes (JavaBeans), and possibly configuration settings--form what's known as a Web application.

Enterprise JavaBeans (EJBs), the other type of J2EE business component that is programmable by developers, is a bit of a different beast. It does not run in a Web Container, but instead requires hosting in an EJB Container. Without going into all the myriad differences between them and their Web application counterparts, suffice it to say that EJBs are oriented toward managing database transactions and toward reading from and writing to relational databases. Applications that combine Web components (JSPs and servlets) with EJBs are known as Enterprise Applications. While you can certainly have an application comprising only JSPs and servlets, EJBs require either a JSP or a servlet to invoke them. Unlike the arrows connecting Web clients to the Web Container, the arrows connecting the Web Container to the EJB Container do not represent HTTP protocols. Instead, a different, more complex, session-oriented wire protocol named IIOP is used (Yes, this is the same IIOP transport that the Domino CORBA classes use).

Finally, the blue circle shows the common API services layer that all major J2EE application server products support. Each of the boxes in this layer represents a standardized Java API such as the following (not a complete list, of course):

- JNDI (Java Naming and Directory Interface) is used to locate servers, directories, and other shared components on the network and to interact with them.
- JDBC (Java Database Connectivity) is a Java version of the more familiar ODBC. It is used to submit queries and read results from relational databases.
- JMS (Java Messaging Service) is another API that is used to implement reliable message queues between components or processes. IBM's MQ, for example, has a JMS interface to it.

You may have noticed that the service layer in the Web Container is identical to that in the EJB Container. This is no accident; both sets of components can make use of the identical services to get their work done.

To create your business application, you write JSPs, servlets, and/or EJBs as appropriate (or get any of a number of tools, such as IBM's Rational Application Developer, to help you generate these components). You then deploy your components in a package to the application server. Once there, it responds to URL invocations (JSPs and/or servlets) or to IIOP requests (EJBs) to perform work.

The WebSphere Application Server platform allows these components to perform well (though, of course, some of the burden is on the developer to write efficient code) by supplying lots of application services (database connections, directory services, message queues, and so on). This relieves the application developer of having to supply code for these services himself and makes application components more portable to other implementations of J2EE because all these services are standardized.

Additionally, WebSphere Application Server is architected so that your application can also be scalable. What does this really mean in practice? If our application (plus the platform on which it runs) is sensitive to external load (that is, more users), we can unblock CPU bottlenecks by distributing pieces of our application across multiple computers without recoding any of the components.This is perhaps the biggest value of an architecture like J2EE; that is, you can, for example, put your JSPs and servlets on one computer and your EJBs on

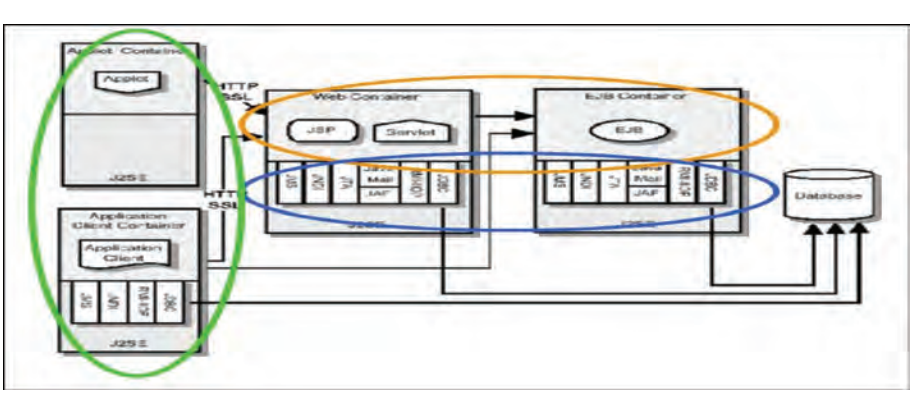

Figure 1. Diagram of basic J2EE architecture, showing important components

another. The connectivity is all managed by WebSphere Application Server; your code does not have to change at all. Thus, tuning for scalability (again, assuming an appropriately coded application) becomes an administrator task, not (primarily) a developer task. Of course, using too many computers for an application will impose a network overhead load, so you must be careful, and there is no one-size-fits-all answer.

This brief summary of what IBM WebSphere Application Server is and how it provides application benefits in terms of performance and scalability is important because WebSphere Application Server is the foundation layer of both IBM WebSphere Portal and IBM Workplace Collaboration Services as we shall see.

What does this mean for Lotus Domino? This is a complicated question, and we have written (or co-written) many articles on various aspects of it. The short (and in some ways inadequate) big-picture answer is that correctly combining the rich collaborative application functionality of Lotus Domino with the speed and scalability of WebSphere Application Server can be a big win. Lotus Domino, for all its great power and richness in workflow, programmability, data integration, and rapid application development, sometimes (for various reasons) just cannot handle exceptionally large application loads. Web agents, in particular, are vulnerable to very high levels of activity because each URL invocation causes the agent to be loaded anew into its target language interpreter (LotusScript, Java, @function, and so on) from the NSF. By contrast, JSPs and servlets in WebSphere Application Server are loaded once, and then the code remains resident in the Web Container's Java Virtual Machine. Concurrent requests are simply invoked on separate Java threads.

The implication is that should you have a situation in which your Domino Web application (which likely is primarily using Web Query Open and Web Query Save agents for its business logic) is not performing well under load you may benefit by recoding some of your Web Query Open/Save agents as servlets running on WebSphere Application Server. Of course, you don't lose important Domino functionality by doing this; you can just continue to use the Domino Java classes.

Bottom line: Integrating Domino Web applications with WebSphere Application Server makes sense if you have performance problems. When we get to WebSphere Portal and Workplace Collaboration Services, you'll see that integrating Lotus Domino with these products solves very different problems.

WebSphere Portal layer: Aggregate UI and device independence

Whereas WebSphere Application Server is all about performance and scalability, WebSphere Portal, in contrast, is all about user experience (what we used to call UI). The major value-adds of WebSphere Portal are that it:

- Brings many application user interfaces (UIs) together in one browser window. This is referred to as UI aggregation.
- Allows personalization in that each Portal user can customize his own screen layout, how much space each application uses on the screen, and where they sit.
- Provides a single sign-on capability, so that (once configured) the user signs on only one time to the Portal, and the Portal figures out how to log into each individual back-end application.
- Provides a large degree of client device independence. Any given application can be exposed both for a PC/browser client and for a hand-held device such as a Palm Pilot without needing special rendering code in the application itself.
- Enables applications co-resident in the Portal to see each other and to exchange data (to a somewhat limited extent due to security concerns). This is often referred to as "on-the-glass integration."

IBM WebSphere Portal is, in fact, a WebSphere Application Server application. It's a collection of JSPs, servlets, and EJBs that is deployed on your WebSphere Application Server configuration (other Portal products work similarly on their respective application server platforms). The Portal code allows developers to implement (or deploy off the shelf) an individual portlet for each backend application type. Portlets are actually just specialized servlets, and their job is to access a particular backend application (a Domino database, a Web site, Domino Web Mail, whatever it may be) and to provide a rendering of that application to WebSphere Portal. The Portal then aggregates all such renderings into the multi-pane screen that the user sees. Figure 2 shows an example of a Portal screen in a browser with some example portlets.

Where do portlets come from? Many are available for free (some are not free) from the IBM Workplace Soluations catalog. If you don't see what you need there, your next easiest step is to use one of the portlet builder tools available from IBM as well as other vendors. These tools provide a no-code interface that lets you describe what you

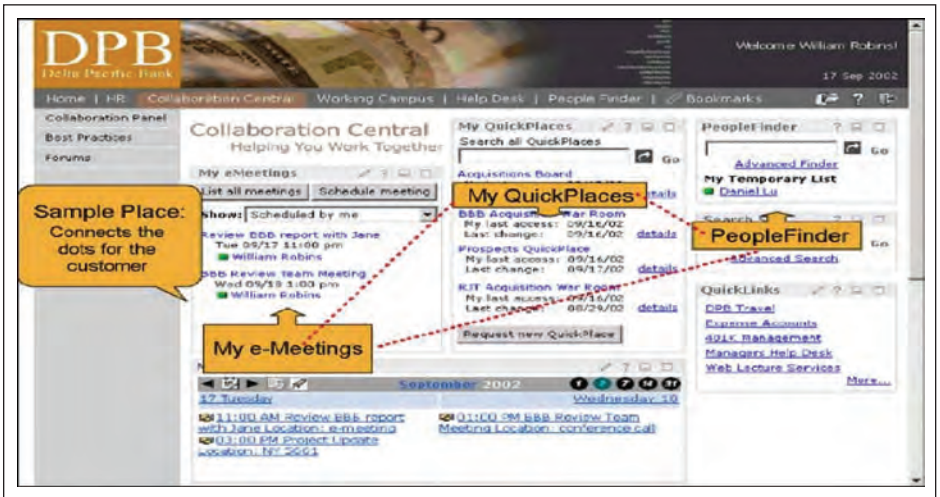

Figure 2. Screenshot of a sample Portal, showing several applications aggregated on the screen

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.

SYS-CON MEDIA

ONLY $69.99

ONE YEAR
12 ISSUES

Subscription Price Includes FREE JDJ Digital Edition!

www.JDJ.SYS-CON.com

or **1-888-303-5282**

OFFER SUBJECT TO CHANGE WITHOUT NOTICE

want your portlet to do, and then it builds the code for you.

If this still doesn't meet your needs, you can use IBM Rational Application Developer, IBM's Eclipse-based development environment, to write your own portlet. Rational Application Developer comes with a Portal test environment and a lot of built-in tools and wizards to help ease the task of portlet building.

The nice thing about the portlet model is that each application can still pretend that it owns all the screen real estate. The job of the portlet is to provide a hook to the backend application and deliver a rendering (typically HTML or XML) of that application to WebSphere Portal, which then figures out where to put it on the screen.

Figure 3 shows an extension of the J2EE diagram depicted in Figure 1 with the Portal layer on top (actually on the side, in this case) and with, as one example, a Domino Web Access portlet accessing a Domino server to provide a rendering of the user's email. The WebSphere Application Server / J2EE diagram has been shrunk down and turned sideways with a Portal layer between it and the client laptop.

Once again, what does this mean for Lotus Domino? WebSphere Portal comes with several in-the-box portlets for exposing Domino NSFs to the Portal interface. Plus, you can use the IBM Portlet Builder for Domino to create and deploy your own wrappers for your Domino applications (or buy one of a number of third-party portlet builders).

Integrating Lotus Domino with WebSphere Application Server can solve performance and scalability issues, whereas integrating Domino applications with WebSphere Portal is a UI thing. You do it to bring your Domino applications into the Portal environment and to reap the benefits of UI aggregation and device independence.

This is an important distinction because doing something for the wrong reason only causes trouble. You wouldn't integrate a Domino application with WebSphere Portal to solve performance problems just as you wouldn't integrate a Domino application with WebSphere Application Server to gain a portal interface.

Workplace layer: Packaged applications and rich client

How do you get from Portal to IBM Workplace Collaboration Services? The answer is by layeringâ . The easiest way to explain it is to say that Workplace Collaboration Services (and related Workplace products, such as Workplace Services Express) is a set of applications implemented on top of WebSphere Portal. In many ways, Workplace Collaboration Services is a Portal application. However, whereas WebSphere Portal is oriented toward browser and hand-held clients, Workplace Collaboration Services also includes a rich client installed (or provisioned) from the server part of the product to users' desktops. This rich client, called Workplace Managed Client , also plays an important role in combination with the next major release of IBM Lotus Notes, code-named Hannover, as we shall see below.

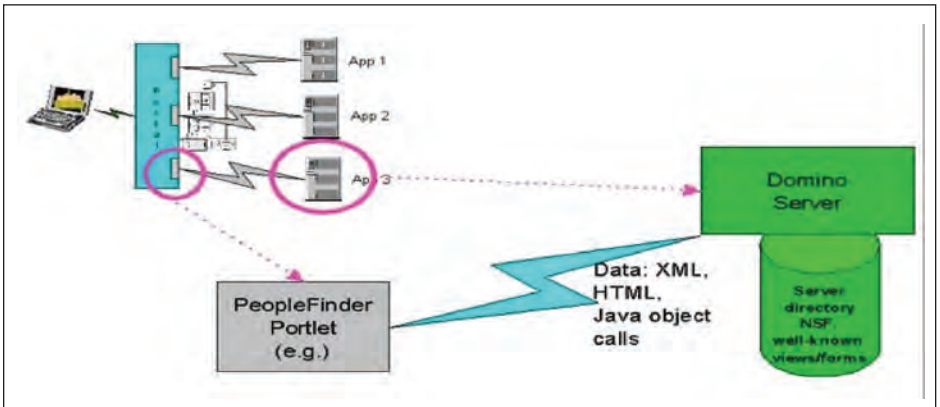

Figure 3. Diagram of WebSphere Portal accessing multiple backend applications through portlets

The key packaged applications that Workplace Collaboration Services delivers on top of the Portal framework include (this list is not meant to be exhaustive):

- Activity Explorer
- Workplace Document Management
- Workplace Web Content Management
- Instant messaging and Web conferencing (essentially a re-hosting of Lotus Sametime on the Workplace framework)
- Team Spaces
- Email (The Workplace Messaging email product is not meant to replace or compete with Notes mail; it is meant for a different audience. Plus, any portlet that can run in WebSphere Portal will also run in Workplace Collaboration Services.)

Both Workplace Collaboration Services and Workplace Services Express are implemented as WebSphere Portal applications, and you can integrate Lotus Domino "on the glass" with both in just the same way you would do it with WebSphere Portal; that is, build, buy, or get off-the-shelf portlets to wrapper your existing Domino applications (as in the example in Figure 3 above).

Workplace Managed Client

One big difference between the two Workplace products is that Workplace Services Express is browser (and hand-held) oriented, while Workplace Collaboration Services also gives you the option of deploying IBM Workplace Managed Client. Users of Lotus Notes already know about rich clients and some of the benefits they provide. Until Workplace Collaboration Services and Workplace Managed Client came along, though, the J2EE platform had no rich client on which to deploy applications.

Workplace Managed Client is quite different from Lotus Notes in many ways, for example:

- Workplace Managed Client is based on the Eclipse open-source framework. Eclipse was originally a developer environment, but it has been put to new use as a client framework. Add-on functionality in Eclipse is implemented as one or more plug-ins, Java programs that occupy space on the screen and provide user interfaces to various things. (Sound suspiciously like a Portal?)
- Workplace Managed Client is a managed

Welcome to the Future of Video on the Web!

REGISTER NOW!
www.iTVcon.com

CALL FOR PAPERS NOW OPEN!

LIVE SIMULCAST!
AROUND THE WORLD ON SYS-CON.TV

iTVCON.COM
INTERNET TV CONFERENCE & EXPO 2006

Coming in 2006 to New York City!

"Internet TV is wide open, it is global, and in true 'Web 2.0' spirit it is a direct-to-consumer opportunity!"

For More Information, Call 201-802-3023 or Email itvcon@sys-con.com

List of Topics:

- Advertising Models for Video-on-demand (VOD)
- Internet TV Commercials
- Mastering Adobe Flash Video
- How to Harness Open Media Formats (DVB, etc)
- Multicasting
- Extending Internet TV to Windows CE-based Devices
- Live Polling During Webcasts
- Video Press Releases
- Pay-Per-View
- Screencasting
- Video Search & Search Optimization
- Syndication of Video Assets
- V-Blogs & Videoblogging
- Choosing Your PVR
- Product Placement in Video Content
- UK Perspective: BBC's "Dirac Project"
- Case Study: SuperSun, Hong Kong

Welcome to the Future!

Did you already purchase your ".tv" domain name?
You can't afford not to add Internet TV to your Website in 2006!

2005 was the year of streaming video and the birth of **Internet TV**, the long-awaited convergence of television and the Internet. Now that broadband is available to more than 100 million households worldwide, every corporate Website and every media company must now provide video content to remain competitive, not to mention live and interactive video Webinars and on-demand Webcasts.

20 years ago the advent of desktop publishing tools opened the doors for the creation of some of today's well-known traditional print media companies as well as revolutionized corporate print communications. Today, with maturing digital video production, the advent of fully featured PVRs, and significant advances in streaming video technologies, **Internet TV** is here to stay and grow and will be a critical part of every Website and every business in the years to come.

It will also very rapidly become a huge challenge to network and cable television stations: **Internet TV** is about to change forever the $300BN television industry, too.

The Internet killed most of print media (even though many publishers don't realize it yet), Google killed traditional advertising models, and **Internet TV** will revolutionize television the way we watch it today. You need to be part of this change!

Jeremy Geelan
Conference Chair, iTVCon.com
jeremy@sys-con.com

PRODUCED BY SYS-CON EVENTS

Track 1	Corporate marketing, advertising, product and brand managers
Track 2	Software programmers, developers, Website owners and operators
Track 3	Advertising agencies, advertisers and video content producers
Track 4	Print and online content providers, representatives from traditional media companies, print and online magazine and newspaper publishers, network and cable television business managers

client, in that the applications and features that each user sees in the client UI is controlled (or provisioned) by the server. The Workplace administrator decides which individuals or groups get which packages, applications, or updates in their copy of the Workplace Managed Client. When users log into the Workplace server from their Workplace Managed Client, new code can be downloaded automatically. This provisioning feature vastly simplifies deployment of client-side code.

- As hinted above, the Workplace Managed Client user experience is quite similar to that of WebSphere Portal except that it's not just a browser. The Workplace Managed Client UI contains an embedded browser, of course, but it also contains a series of productivity editors or tools, including a word processor, a spreadsheet, and a presentation / slide show editor. These editors provide most of the word processing and spreadsheet functionality most people need on a day-to-day basis. They can read Microsoft Office files, but instead of saving data in proprietary file formats, the Workplace editors use XML.

The Eclipse / Workplace Managed Client plug-in architecture is quite powerful. In fact, one of the plug-ins you can use with Workplace Managed Client is Lotus Notes 7 (you must install Lotus Notes 7 independently on the computer, but if it's there, then Workplace Managed Client sees it and allows you to launch Lotus Notes in the Workplace Managed Client window).

The Workplace Managed Client thus becomes much more than just a browser window with a bunch of applications in it. It lays the foundation for true composite applications by taking integration on the glass to the next level. Workplace Managed Client is a programmable client -- you can write your own Eclipse plug-ins, and you can use the Workplace Managed Client APIs distributed by IBM to enable client-side components (whether they be traditional portlets, Notes-based applications, or custom-coded plug-ins) to communicate with each other and share data (with appropriate security, of course).

This kind of client-side functionality both incorporates Lotus Notes (with full support for replication and off-line use) and goes beyond it. The Eclipse UI is customizable and extensible in ways that Lotus Notes never was. Would you like a custom UI for a calendar view? No problem, write your own plug-in viewer and hook it into the Workplace Managed Client event model. Would you like something special to happen when users send email? No problem, write a custom event handler. (Of course, you need to be a Java programmer). What does this mean for the future of the traditional Notes client? Is it going away? Absolutely not. Is IBM Workplace replacing Lotus Domino? Absolutely not. In fact, the two platforms are evolving and merging in important ways that take nothing away from the Notes/Domino functionality we all know and love.

Evolution of the Notes and Workplace rich clients: Merger of equals

Today, the Notes client and the Workplace Managed Client look like two vastly different things. This situation naturally leads to all sorts of speculation about which one will "win" and whether or not people will be forced to migrate applications from one platform to another.

The good news is that this situation is only temporary. The two client platforms (Lotus Notes and Workplace Managed Client) will actually be merging together in the next major release of the traditional Notes client, code-named Hannover, after the city in Germany where the product was first announced. Hannover is truly innovative: It's still Lotus Notes underneath, but the entire UI has been re-hosted on the Eclipse open framework, making Notes components first-class players in the composite application client space. Furthermore, Notes components will be provisioned from either a Domino server or a Workplace server.

What this really means is that by upgrading to Hannover, customers will also gain the benefits of the Workplace rich client--including portal interface, server-management/provisioning, extensible UI--while it's still Lotus Notes. Moreover, because it's still a real Notes client underneath the Eclipse wrappings, all your existing Notes/Domino applications will continue to run unmodified in the Hannover environment.

Of course, traditional stand-alone Notes and Domino configurations will continue to be offered by IBM. No one will be forced to adopt the new Hannover client against their will. The benefits of the unified rich client, though, are huge, and we predict that most customers will find the transition smooth and the results very empowering.

Different layers, different benefits

Why IBM Workplace then? It's worth repeating (if only because there has been so much confusion about this over the past couple of years) that IBM Lotus Notes and Domino are not going away.

The J2EE foundation technology embodied in IBM WebSphere Application Server provides a fast, scalable platform onto which organizations can deploy applications. By integrating Lotus Domino with WebSphere Application Server, you can unlock some of the performance bottlenecks inherent in high-volume Domino Web applications without losing any of the rich collaborative functionality that Lotus Domino provides.

The Portal layer on top of WebSphere Application Server provides personalized desktops on which knowledge workers can organize their workspace as they choose and bring all their important applications into one place with integration on the glass of their data.

The IBM Workplace layer on top of WebSphere Portal gives you a series of prepackaged applications, the ability to expose existing Domino applications in a Portalized environment, and powerful new administrative options for provisioning. Furthermore, Workplace Collaboration Services now provides users with the Workplace Managed Client that further enhances the on-line work experience. The recently shipped version of Workplace Managed Client supports a plug-in for Lotus Notes 7. The next major release of Lotus Notes will actually be a merger of the traditional Notes client with Workplace Managed Client, yielding powerful enhancements of both. This is not at all a replacement of Lotus Notes with Workplace Managed Client rather it's a merger of equals that make both better.

Part One of a Four-Part Article

Demystifying Class Loading Problems

BY LAKSHMI SHANKAR AND SIMON BURNS

The class loading component is fundamental to the Java virtual machine. Though developers generally have a good grasp of the basics of class loading, when a problem occurs, they may have a hard time diagnosing and determining a solution. In this four-part article series, we discuss the various class loading problems that you may encounter in your Java development and illustrate why they occur and how to resolve them. These insights should help you understand and resolve common Java exceptions, such as NoClassDefFoundError and ClassNotFoundException, in addition to more challenging problems, such as class loader constraint violations and deadlocks. In this first article, we describe in detail how Java class loading works and discuss the tools available in the JVM to help you diagnose class loading problems.

Class loaders are responsible for loading classes into the Java Virtual Machine (JVM). Simple applications can use the Java platform's built-in class loading facility to load their classes; more complex applications tend to define their own custom class loaders. No matter what kind of class loader you're using, however, there are many problems that can occur in the class loading process. If you want to avoid such problems, you need to understand the fundamental mechanics of class loading. When problems do occur, an appreciation of the available diagnostic features and debugging techniques should help you resolve them.

In this series of articles, we'll provide an in-depth look at class loading problems and use comprehensive examples to illustrate them. The first section of this introductory article describes the fundamentals of class loading; the second introduces some JVM debugging features. The next three articles in this series will focus on resolving class loading exceptions and illustrate some of the trickier class loading problems that you may encounter.

Fundamentals of class loading

This section describes the core concepts of class loading to provide a knowledge base for the rest of this series.

Class loader delegation

The *class loader delegation model* is the graph of class loaders that pass loading requests to each other. The *bootstrap* class loader is at the root of this graph. Class loaders are created with a single *delegation parent* and looks for a class in the following places:

- Cache
- Parent
- Self

A class loader first determines if it has been asked to load this same class in the past. If so, it returns the same class it returned last time (that is, the class stored in the cache). If not, it gives its parent a chance to load the class. These two steps repeat recursively and depth first. If the parent returns null (or throws a ClassNotFoundException), then the class loader searches its own path for the source of the class.

Because the parent class loader is always given the opportunity to load a class first, the class is loaded by the class loader nearest the root. This means that all core bootstrap classes are loaded by the bootstrap loader, which makes sure that the correct versions of classes such as java.lang.Object are loaded. This also has the effect of only allowing a class loader to see classes loaded by itself or its parent or ancestors; it cannot see classes loaded by its children.

(This article was first published on developerWorks WebSphere at http://www.ibm.com/developerWorks/websphere.)

ABOUT THE AUTHORS

Lakshmi Shankar is a Software Engineer in IBM Hursley Labs, UK. He has worked for IBM for more than three years and has a broad range of experience, having worked in Java performance, test, and development within Hursley Labs. Until recently he was the Class Loading component owner for IBM's Java technology. He is now a developer working as part of the Information Management team.

Simon Burns was the component owner and team lead for the Persistant Reusable JVM in the Java Technology Centre in IBM Hursley Labs. He worked in JVM development for over three years, specializing in the Persistant Reusable JVM technology and the z/OS platform. He has also worked closely with CICS, helping them to exploit this technology. Simon worked on the OSGi framework as part of the open-source Eclipse Equinox project, which has now been integrated into Eclipse 3.1. He is now working on componentization.

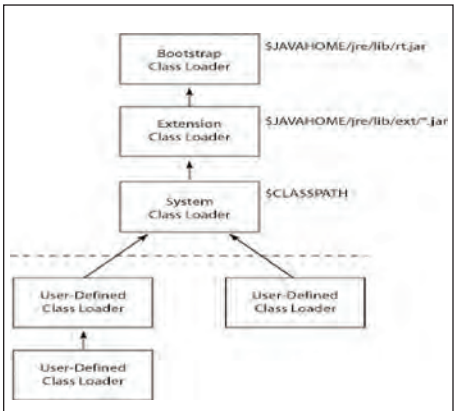

Figure 1. The class loader delegation model

Figure 1 shows the three standard class loaders.

Unlike all other class loaders , the bootstrap class loader (also known as the *primordial* class loader) cannot be instantiated by Java code. (Often, this is because it is implemented natively as part of the VM itself.) This class loader loads the core system classes from the boot classpath, which is normally the JAR files located in the jre/lib directory. However, you can modify this classpath using the -Xbootclasspath command-line options (which we describe later).

The *extension* class loader (also known as the *standard extensions* class loader) is a child of the bootstrap class loader. Its primary responsibility is to load classes from the extension directories, normally located the jre/lib/ext directory. This provides the ability to simply drop in new extensions, such as various security extensions, without requiring modification to the user's classpath.

The *system* class loader (also known as the *application* class loader) is the class loader responsible for loading code from the path specified by the CLASSPATH environment variable. By default, this class loader is the parent of any class loader created by the user. This is also the class loader returned by the ClassLoader.getSystemClassLoader() method.

Classpath options

Table 1 summarizes the command-line options for setting the classpaths of the three standard class loaders:

The phases of class loading

The loading of a class can essentially be broken down into three phases: loading, linking, and initializing.

Most, if not all, problems relating to class loading can be tracked down to a problem occurring in one of these phases. Therefore, a thorough understanding of each phase helps in the diagnosing of class loading problems. The phases are illustrated in Figure 2:

The *loading* phase consists of locating the required class file (by searching though the respective classpaths) and loading in the bytecode. Within the JVM, the loading process gives a very basic memory structure to the class object. Methods, fields, and other referenced classes are not dealt with at this stage. As a result, the class is not usable.

Linking is the most complicated of the three phases. It can be broken down into three main stages:

- ***Bytecode verification***. The class loader does a number of checks on the bytecodes of the class to ensure that it is well formed and well behaved.
- ***Class preparation***. This stage prepares the necessary data structures that represent fields, methods, and implemented interfaces that are defined within each class.
- ***Resolving***. In this stage, the class loader loads all the other classes referenced by a particular class. The classes can be referenced in a number of ways:
  - Superclasses
  - Interfaces
  - Fields
  - Method signatures
  - Local variables used in methods

During the *initializing* phase, any static initializers contained within a class are executed. At the end of this phase, static fields are initialized to their default values.

At the end of these three phases, a class is fully loaded and is ready for use. Note that class loading can be performed in a lazy manner and therefore some parts of the class loading process may be done on first use of the class rather than at load time.

Explicit vs. implicit loading

There are two ways in which classes can be loaded -- *explicitly* or *implicitly* -- with subtle variations between the two.
Explicit class loading occurs when a class is loaded using one of the following method calls:

- cl.loadClass() (where cl is an instance of java.lang.ClassLoader)
- Class.forName() (the starting class loader is the defining class loader of the current class)

When one of these methods is invoked, the class whose name is specified as an

Command-line option	Explanation	Class loader involved
-Xbootclasspath:<directories and zip/JAR files separated by ; or :>	Sets the search path for bootstrap classes and resources.	Bootstrap
-Xbootclasspath/a:<directories and zip/JAR files separated by ; or :>	Appends the path to end of the boot classpath.	Bootstrap
-Xbootclasspath/p:<directories and zip/JAR files separated by ; or :>	Prepends the path to front of the boot classpath.	Bootstrap
-Dibm.jvm.bootclasspath=<directories and zip/JAR files separated by ; or :>	The value of this property is used as an additional search path, which is inserted between any value that is defined by -Xbootclasspath/p: and the boot classpath. The boot classpath is either the default, or the path that is defined by the -Xbootclasspath: option.	Bootstrap
-Djava.ext.dirs=<directories and zip/JAR files separated by ; or :>	Specifies the search path for the extension classes and resources.	Extension
-cp or -classpath <directories and zip/JAR files separated by ; or :>	Sets the search path for application classes and resources.	System -Djava.class.
path=<directories and zip/JAR files separated by ; or :>	Sets the search path for application classes and resources.	System

Table 1. Classpath options

argument is loaded by the class loader. If the class is already loaded, then a reference is simply returned; otherwise, the loader goes through the delegation model to load the class.

Implicit class loading occurs when a class is loaded as result of a reference, instantiation, or inheritance (not via an explicit method call). In each of these cases, the loading is initiated under the covers and the JVM resolves the necessary references and loads the class. As with explicit class loading, if the class is already loaded, then a reference is simply returned; otherwise, the loader goes through the delegation model to load the class.

Classes are often loaded through a combination of explicit and implicit class loading. For example, a class loader could first load a class explicitly and then load all of its referenced classes implicitly.

Debugging features of the JVM

The previous section introduced the fundamental principles of class loading. This section covers the variety of features built into the IBM JVM to assist debugging. Other JVMs may have similar debugging features available; refer to the appropriate documentation for details.

Verbose outputYou can turn on the IBM JVM's verbose output by using the -verbose command-line option. Verbose output displays information on the console when certain events occur -- when a class has been loaded, for instance. For additional class loading information, you can use verbose class output. This is activated using the -verbose:class option.

Interpreting verbose output

Verbose output lists all the JAR files that have been opened and includes the full path to those JARs. Here's an example:

```
...
[Opened D:\jre\lib\core.jar in 10 ms]
[Opened D:\jre\lib\graphics.jar in 10 ms]
...
```

All the classes that are loaded are listed, along with the JAR file or directory from which they were loaded. For example:

```
...
[Loaded java.lang.NoClassDefFoundError from
D:\jre\lib\core.jar]
[Loaded java.lang.Class from D:\jre\lib\core.
jar]
[Loaded java.lang.Object from D:\jre\lib\core.
jar]
...
```

Verbose class output shows additional information, such as when superclasses are being loaded, and when static initializers are being run. Some example output follows:

```
...
[Loaded HelloWorld from file:/C:/myclasses/]
[Loading superclass and interfaces of HelloWorld]
[Loaded HelloWorldInterface from file:/C:/myclass-
es/]
[Loading superclass and interfaces of
HelloWorldInterface]
[Preparing HelloWorldInterface]
[Preparing HelloWorld]
[Initializing HelloWorld]
[Running static initializer for HelloWorld]
...
```

Verbose output also shows some internally thrown exceptions (if they occur), including the stack trace.

Resolving problems using -verbose

Verbose output can help to solve classpath problems, such as JAR files not being opened (and therefore not on the classpath) and classes being loaded from the wrong place.

IBM Verbose Class LoadingIt is often useful to know where class loaders look for classes and which class loader loads a particular class. You can obtain this information using the IBM Verbose Class Loading command-line option: -Dibm.cl.verbose=<class name>. You can use regular expressions to declare the name of the class; for instance, Hello* traces any classes with names starting with Hello.

This option also works on user-defined class loaders, as long as they directly or indirectly extend java.net.URLClassLoader.

Interpreting IBM Verbose Class Loading output

IBM Verbose Class Loading output shows the class loaders that attempt to load the specified class and the locations in which they look. For example, imagine we used the following command line:

```
java -Dibm.cl.verbose=ClassToTrace MainClass
```

Here, MainClass references ClassToTrace in its main method. This would produce output similar to the output here.

The class loaders are listed with parents before children because of the way that the standard delegation model works: Parents go first.

Notice that there is no output for the bootstrap class loader. Output is only produced for class loaders that extend java.net.URLClassLoader. Note also that class loaders are listed under their class name; if

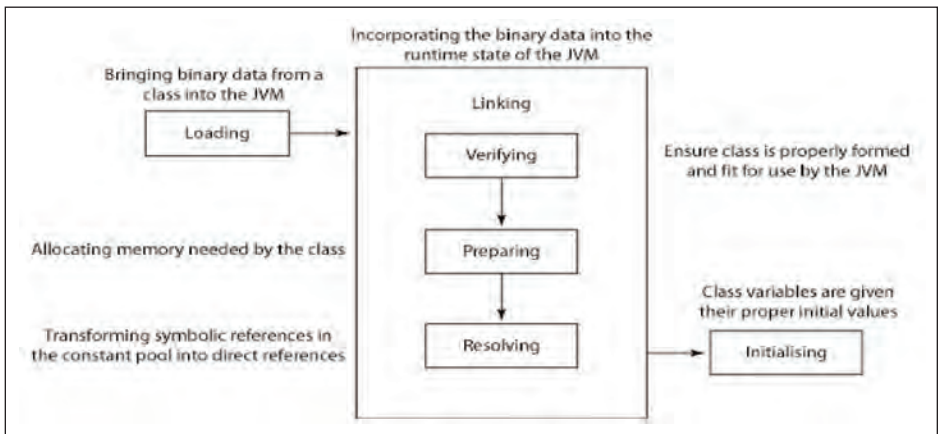

Figure 2. The phases of class loading

Subscribe Today!

– INCLUDES –
FREE DIGITAL EDITION!
(WITH PAID SUBSCRIPTION)
GET YOUR ACCESS CODE INSTANTLY!

The major infosecurity issues of the day... ***identity theft, cyber-terrorism, encryption, perimeter defense, and more*** *come to the forefront in* **ISSJ** *the storage and security magazine targeted at IT professionals, managers, and decision makers*

Editorial Mission

Greater Collaboration and Efficiency Through Education

✓ ***ISSJ*****'s editorial mission is to showcase proven solutions** that will guide, motivate, and inspire senior IT and business management leaders in the planning, development, deployment, and management of successful enterprise-wide security and storage solutions.

✓ ***ISSJ*** brings together all key elements of data storage and protection, and presents compelling insight into the benefits, efficiencies, and effectiveness gained by focusing on these two critical areas of IT simultaneously.

✓ ***ISSJ*** is an objective, critical source of information that helps storage and security managers make informed management decisions about what is working today, and what they need to plan for tomorrow, and is the only publication that focuses exclusively on the needs of IT professionals who are driving the enterprise storage architecture/ infrastructure while pursuing and incorporating the latest security technologies.

✓ ***ISSJ*** achieves our mission by delivering in-depth features, practical "how-to" solutions, authoritative commentary, hard-hitting product reviews, comprehensive real-world case studies, and successful business models, written by and for professional IT storage and security practitioners.

SAVE 50% OFF!
(REGULAR NEWSSTAND PRICE)

Only $39.99
ONE YEAR
12 ISSUES

www.ISSJournal.com
or
1-888-303-5282

SYS-CON MEDIA
The World's Leading i-Technology Publisher

there are two instances of a class loader, it may not be possible to distinguish between them.

Resolving problems using IBM Verbose Class Loading

The IBM Verbose Class Loading option is a great way to check what the classpaths for all class loaders have been set to. It also indicates which class loader loads a given class and where it loads it from. This makes it easy to see if the correct version of a class is being loaded.

JavadumpA Javadump (also know as a Javacore) is another IBM diagnosis tool that you may find useful; to learn more about it, see the IBM Diagnostics Guide (see Resources for a link). A Javadump is generated by the JVM when one of the following events occurs:

- A fatal native exception occurs
- The JVM runs out of heap space
- A signal is sent to the JVM (for example, if Control-Break is pressed on Windows, or Control-\ on Linux)
- The com.ibm.jvm.Dump.JavaDump() method is called

The moment that Javadump is triggered, detailed information is recorded in a date-stamped text file saved in the current working directory. This information includes data about threads, locks, stacks, and so on, as well as a rich set of information about the class loaders in the system.

Interpreting the class loading section of a Javadump

The exact information that is provided in a Javadump file depends on the platform on which the JVM is running. The class loader section includes:

- The defined class loaders and the relationship between them
- Lists of classes loaded by each class loader

The following is a snapshot of the class loader information taken from a Javadump:

```
CL subcomponent dump routine
============================
 Classpath 2(D:\jre\lib\core.jar),...
 Oldjava mode false
 Bootstrapping false
 Verbose class dependencies false
 Class verification VERIFY_REMOTE
 Namespace to classloader 0x00000000
 Start of cache entry pool 0x44D85430
 Start of free cache entries 0x44D86204
 Location of method table 0x44C23AA0
 Global namespace anchor 0x00266894
 System classloader shadow 0x00376068
 Classloader shadows 0x44D7BA60
 Extension loader 0x00ADB830
 System classloader 0x00ADB7B0
 Classloader summaries
         12345678: 1=primordial,2=extension,3=shar
eable,4=middleware,
                   5=system,6=trusted,7=applicat
ion,8=delegating
         -----ta- Loader sun/misc/Launcher$AppClas
sLoader(0x44D7BA60),
                  Shadow 0x00ADB7B0,
                  Parent sun/misc/Launcher$ExtCla
ssLoader(0x00ADB830)
                  Number of loaded classes 1
                  Number of cached classes 260
                  Allocation used for loaded
classes 1
                  Package owner 0x00ADB7B0
         -xh-st-- Loader sun/misc/Launcher$ExtClas
sLoader(0x44D71288),
                  Shadow 0x00ADB830,
                  Parent *none*(0x00000000)
                  Number of loaded classes 0
                  Number of cached classes 0
                  Allocation used for loaded
classes 3
                  Package owner 0x00ADB830
         p-h-st-- Loader *System*(0x00376068),
Shadow 0x00000000
                  Number of loaded classes 304
                  Number of cached classes 304
                  Allocation used for loaded
classes 3
                  Package owner 0x00000000
 ClassLoader loaded classes
         Loader sun/misc/Launcher$AppClassLoader(
0x44D7BA60)
                  HelloWorld(0x00ACF0E0)
         Loader sun/misc/Launcher$ExtClassLoader(
0x44D71288)
         Loader *System*(0x00376068)
                  java/io/WinNTFileSystem(0x002C
D118)
                  java/lang/Throwable(0x002C03A
8)
                  java/lang/IndexOutOfBoundsExcep
tion(0x44D45208)
                  java/lang/UnsatisfiedLinkError(
0x44D42D38)
...................classes left out to save space
.....................
                  [Ljava/lang/Class;(0x002CA9E8)
                  java/io/InputStream(0x002C981
8)
                  java/lang/Integer$1(0x002C83E
8)
                  java/util/Dictionary(0x002C4298)
```

In this example, there are only the three standard class loaders:

- The system class loader (sun/misc/

"The class loader delegation model is the graph of class loaders that pass loading requests to each other"

Launcher$AppClass loader)
- The extension class loader (sun/misc/Launcher$ExtClass loader)
- The bootstrap class loader (*System*)

The Classloader summaries section provides details about each class loader in the system, including the type of the class loader. In this series of articles, the types of interest are the *primordial, extension, system, application,* and *delegation* (used in reflection). The other types (*shareable, middleware,* and *trusted*) are used in the Persistent Reusable JVM, which is beyond the scope of these articles (see the Persistent Reusable JVM User Guide for more information; there's a link in the Resources section below). The summaries section also shows the parent class loader: The parent of the system class loader is sun/misc/Launcher$ExtClass loader(0x00ADB830). This parent address corresponds to the native data structure of the parent class loader (called the *shadow*).

The ClassLoader loaded classes section lists the classes loaded by each class loader. In this example, the system class loader has only loaded one class, HelloWorld (at address 0x00ACF0E0).

Resolving problems using Javadumps

Using the information provided in the Javadump, it is possible to ascertain which class loaders exist within the system. This includes any user-defined class loaders. From the lists of loaded classes, it is possible to find out which class loader loaded a particular class. If the class cannot be found, that means that it was not loaded by any of the class loaders present in the system (which would usually result in a ClassNotFoundException).

Other types of problems that could be diagnosed using a Javadump include:
- ***Class loader namespace problems***. A class loader namespace is a combination of a class loader and all the classes that it has loaded. For example, if a particular class is present but is loaded by the wrong class loader (sometimes resulting in a NoClassDefFoundError), then the namespace is incorrect -- that is, the class is on the wrong classpath. To rectify such problems, try putting the class in a different location -- in the normal Java classpath, for instance -- and make sure that it gets loaded by the system class loader.
- ***Class loader constraint problems***. We'll discuss an example of this kind of problem in the final article in this series.

Java method tracingThe IBM JVM has a built-in method tracing facility. This allows methods in any Java code (including the core system classes) to be traced without modification to the code. Because this facility can provide a large amount of data, you can control the level of trace in order to zero in on the information that you want.

The option for enabling trace varies depending on the release of the JVM. For details of these options, refer to the IBM Diagnostics Guides.

Here are some example command lines:

To trace all java.lang.ClassLoader methods when running HelloWorld in IBM Java 1.4.2:

```
java -Dibm.dg.trc.methods=java/lang/
ClassLoader.*() -Dibm.dg.trc.print=mt HelloWorld
```

To trace the loadClass() method in ClassLoader and the methods in HelloWorld, also in IBM Java 1.4.2:

```
java -Dibm.dg.trc.methods=java/lang/ClassLoader.
loadClass(),HelloWorld.*()
  -Dibm.dg.trc.print=mt HelloWorld
```

Interpreting method trace output

Here is a sample of method trace output (using the second command line from the previous section).

Each line of trace provides more information than shown above. Let's look at one of the above lines in full:

```
12:57:41.277 0x002A23C0 04000D   > java/lang/
ClassLoader.loadClass Bytecode method,
       This = 0x00D2B830, Signature: (Ljava/
lang/String;Z)Ljava
/lang/Class;
```

This tracing includes:
- 12:57:41.277: The timestamp of method entry or exit.
- 0x002A23C0: The thread ID.
- 04000D: An internal JVM trace point used by some advanced diagnostics.
- The remaining information shows whether a method is being entered (>) or exited (<), followed by details of the method.

Resolving problems with method trace

Method tracing can be used to resolve different types of problems, including:
- Performance hotspots: Using timestamps, it is possible to find methods that take a significant amount of time to execute.
- Hangs: The last method entry is usual a good indication of where the application has hung.
- Incorrect objects: Using the address, it is possible to check that methods are being invoked on the desired object by matching to the address on the constructor call for that object.
- Unexpected code paths: By following the entry and exit points, it is possible to see any unexpected code paths taken by the program.
- Other faults: The last method entry is usual a good indication of where the fault has occurred.

What's next

In this article, you learned the fundamentals of class loading in JVMs and the debugging features available in the IBM JVM. In the next article in this series, you'll learn how to apply this knowledge to understand and resolve various class loading problems typically encountered when running Java applications.

NEWS ROUND-UP

IBM Steps Up Effort to Drive Sotware as a Service

IBM has announced new resources to help application providers develop and deploy business solutions to be delivered as services. IBM will provide companies with access to technical architects, free online educational courses and new sales incentives to help them transform their applications and business operations to a services delivery model. As a result, application providers can now more quickly deploy software as a service offerings and increase profits. In an effort to continue to help business partners capture growth in this booming market, IBM will provide new sales and marketing incentives to business partners. Starting immediately, channel partners can now receive a 10 percent referral fee in return for submitting a lead to a software as a service application provider that results in closed business. I BM will also provide business partners with access to its sales force to help close opportunities faster and gain credibility with IBM's sales experts locally.

Additionally, IBM is announcing new strategic relationships with Apptix, Success Factors, and Corente who are delivering their applications as services based on IBM's open technology infrastructure and support from IBM services, which provides a secure hosting environment for 24/7 network reliability and availability. In less than one year IBM has doubled its channel network delivering software as a service to more than 100 partners. Working with IBM, these companies have reached tens of thousands of customers and millions of users in a services model.

Pacific Crest estimates the market for hosted software will grow 25% a year to $10 billion by 2009. AMR Research's 2005 Software-as-a-Service syndicated research study shows that more than 78 percent of 500 respondents across major vertical industries and company sizes are currently using or considering software-as-a-service. Only 18 percent said they have no plans to consider software as a service.

Software as a service is a model of delivering software over the Internet, eliminating the need for companies to buy, build, manage and maintain applications, such as accounting, human resources, CRM, ERP, etc. Companies are increasingly drawn to this concept to reduce operational cost, lower maintenance expenses, enable rapid implementation and increase profits.

Today's news builds on IBM's efforts to deliver the broadest range of hardware, software and services resources to help business partners to help business partners quickly deploy this new model of computing. IBM launched its "Software as Services Showcase" initiative less than eight months ago, an online directory which allows customers to search for ISV solutions that are available to them as a service.

For example, San Mateo, CA-based Success Factors, a provider of on-demand, Web-based integrated talent management solutions, joined IBM's Software as Services Showcase and leveraged IBM's direct mail, telemarketing and advertising and technical resources to generate customer leads. In 2005, IBM also sponsored 50 percent regional tradeshows for SuccessFactors' HR applications.

Partnering with IBM has enabled SuccessFactors to generate more than 800 sales leads in less than a year. The Showcase is also helping Success Factors build a channel network -- partners within the Showcase can align their offerings to deliver complementary offerings for customers, and cross sell within the Software as a Service community of IBM's channel partners.

"IBM's support for helping solution providers adopt the Software as a Service model is a significant endorsement for the future of delivering hosted services, a key element of driving technology innovation," said Roger Goulard, vice president of alliances, Success Factors. "Today, more than one million users in 101 countries trust our hosted software solution for talent and workforce management needs, hosted by IBM. Partnering with IBM has allowed us to reach a larger and more varied set of audience than we would have been able to reach on our own. The technical, marketing and sales resources provided by IBM have also allowed us to maximize our reach to newer markets."

New Sales Incentives and Marketing Resources offered to Business Partners

IBM is introducing new ways to invest and promote Business Partners' solutions being delivered as a service. Channel partners can now receive a 10 percent referral fee in return for submitting a lead to a software as a service application provider that results in closed business. Other new resources include IBM hosted web seminars for demand generation to drive customer awareness of software as a service, discounted advertising for ISVs in industry publications and lead generation services and white papers from IBM experts for business insight.

New Technical Workshop - Hybrid Transformation Model

IBM is offering free technical workshops to Business Partners at IBM Innovation Centers worldwide to help ISVs evaluate, adopt and deploy software as a service business model. The workshop is facilitated by IBM enablement architects who will discuss the technical considerations with the solution provider before building and offering software as a service applications model. Architects will also provide advice needed to reach business goals and transition to this model. The workshop includes education and guidance on the architecture of software as a service and new code design methods including Web services, service-oriented architecture (SOA), and security.

New IBM Virtual Innovation Center (VIC) Support

IBM is offering Business Partners worldwide access to new virtual workshops on hardware, software and services support around Software as a Service through the IBM Virtual Innovation Center. Developers can access these workshops on Software as a Service online and complete the courses at their own pace. IBM also offers classes through its virtual mentoring program, allowing participants to attend instructor-led workshops by accessing web conferencing technologies. IBM's new catalogue of Software as a Service classes are designed to help developers transition to a Software as a Service model -- from the early stages of learning about the opportunity to understanding the final implantation procedures. The courseware includes topics such as software as a service transformation techniques, middleware of software as a service, financial implication of e-business on demand, etc.

Expanded Sales Coverage Opportunities through Sales Connections

Through IBM's Sales Connections, ISVs worldwide who are working on active sales opportunities around software as a service can now connect with an IBM local software as a service sales expert to help the ISV close the deal. This is specifically designed to help ISVs gain access to IBM's sales force enabling them to shorten the sales cycle, close business faster and gain instant credibility with new customers with IBM's sales support.

"Making the transition to an Internet delivery model unfolds new market opportunities, but can pose a technical challenge for software makers accustomed to a client server model. It requires major changes underlying software architecture and in sales and marketing strategies, alterations in revenue streams and a cultural change within an organization," said Buell Duncan, general manager, ISV and Developer Relations. "We are dedicated to the success of our business partners, helping them understand the technical and business benefits of transitioning to a software as a service model, find new ways to manage change and support their customers while lowering the costs of doing business to maximize profits."
The Software as Services Showcase can be accessed at: www.ibm.com/software/showcase/sas.

The news is another major step in IBM's commitment to the success of its business partners. As part of its $1 billion investment on joint sales and marketing activities through IBM's PartnerWorld Industry Networks program, IBM has unveiled an array of resources to increase business partner visibility in the marketplace including a discounted advertising program, lead generation tools and direct mail campaign support. More information on PartnerWorld Industry Networks can be accessed at www-304.ibm.com/jct09002c/isv/marketing/industrynetworks/.

IBM's Tivoli Aims to Reduce Clutter

IBM has announced software to help people reduce the clutter and confusion caused by nonstop requests to enter or change their user IDs and passwords. The software -- IBM Tivoli Access Manager for Enterprise Single Sign-on -- broadens IBM's identity and password management capabilities. It consists of technology from Passlogix, which adds single password access for Lotus Notes, SAP and Microsoft Windows-based applications, among others, to IBM's password management support.

By adding this technology to IBM's password management software, people can reach more applications -- from email to vendor extranets -- via "single sign-on," which requires only one password to access a variety of applications. It also enables a greater range of applications to support "two-factor authentication" -- where two types of verification, such as smart cards or biometrics, are needed to log on to a computer.

The software enables people to use computers shared by co-workers to securely access confidential information -- without having to shut down the computer to ensure confidentiality, the company says.

For example, doctors and nurses can sign on to a hospital computer to access medical records and update patient information. Should they forget to log off, the software monitors and protects the workstation by ending inactive sessions and requires the next user to sign on to enter the system.

"People are tired of keeping track of user names and passwords, so they resort to careless tactics like writing passwords on notes stuck to their computer monitors," said Hershel Harris, Vice President of Development, IBM Tivoli Storage and Security. "By teaming with Passlogix, IBM is extending our identity, access and compliance portfolio to help customers cut the clutter out of password management, while helping them satisfy compliance requirements."

The software broadens the capabilities of IBM's Tivoli Access Manager for e-Business and Tivoli Identity Manger, which together with the new software comprise a complete identity management solution. With IBM's identity management software, employees at government agencies, financial institutions and healthcare facilities don't have to remember dozens of passwords and user IDs, each having different character requirements and expiration dates. By improving the way passwords are managed, employees won't incur help desk costs because they forget or lose their passwords.

CA Pledges Day One Support of IBM's System z9

CA has pledged "day one" support, and plans for early exploitation of IBM's newly announced System z9 Integrated Information Processor (zIIP), a specialty engine that will run portions of eligible workloads designed to help free up general computing capacity on the IBM System z9 mainframe.In exploiting zIIP, CA will enable customers to effectively manage and leverage enterprise computing environments that include zIIP resources. CA also will take advantage of zIIP to increase the capabilities and improve the TCO of many of its industry-leading solutions, including its security and systems management software, mainframe DBMS, and DB2 tools

By providing day one support for zIIP, CA will ensure that its customers can immediately take advantage of the new processor's capabilities to cost effectively manage scalability and application performance -- while maintaining rigorous information security and protecting business continuity.

"IBM's zIIP processor is designed to improve the scalability and performance of both core business applications and enterprise IT management tools running on the mainframe," said Russell Artzt, executive vice president of products at CA. "CA is committed to helping our customers get maximum value from their investment in zIIP technology."

CA also will provide tools that help customers leverage zIIP and other technologies to help secure and optimize the performance and scalability of XML, Java and other Web services in the mainframe environment.

"CA has a history of providing support for new technologies from IBM," said Terri Virnig, vice president of IBM System z. "We welcome CA's participation in the introduction of this exciting mainframe enhancement. We expect that zIIP will help the global mainframe community transform their information assets into actionable intelligence to drive business value."

IBM Announces New Version of Workplace Products

IBM has announced availability of new

releases of its core IBM Workplace(TM) products, including Workplace Collaboration Services 2.6, Workplace Managed Client(TM) 2.6, Workplace Forms(TM) 2.6 and Workplace Designer 2.6. The 2.6 versions of these products include more than 50 new features and functions, including enhancements to cross-platform and open standards support. Along with WebSphere Portal, the Workplace family of products help provide the front-end to IBM's Services Oriented Architecture (SOA) strategy.

Support for SOA and open standards, including the OpenDocument Format (ODF) standard in the Workplace Managed Client, iCal support in Workplace Collaboration Services, and XForms in Workplace Forms, gives IBM's Workplace customers greater flexibility and choice. Key innovations such as activity-centric computing provide organizations with tools to transform how employees work together.

SOA enables new applications and business processes to be assembled from reusable components or "services," giving customers the flexibility to adapt to changing business conditions and react to market opportunities. IBM Workplace software products offer a rich set of ready-to-use collaboration components, deployed as SOA services that can be combined with line-of-business applications to create customized, role-based work environments.

New features and functions in the new version of IBM's Workplace family of software products include:

Workplace Collaboration Services 2.6

IBM Workplace Collaboration Services is a single, integrated product based on open standards that includes a wide range of collaborative capabilities or "services," such as e-mail, calendaring, team spaces, instant messaging, on-line learning, Web conferencing, document and Web content management. The 2.6 release includes user interface enhancements across all components, an enhanced document search engine, support for ODF, iCal support for calendar interoperability with IBM Lotus® Notes®, and a new instant messaging gateway to Lotus Sametime, which will connect communities of users between IBM Lotus Sametime and Workplace Collaboration Services. Additions to Workplace Collaboration Services cross-platform support include clustering for iSeries, DB2 Enterprise Server Edition V8.2 as a data store, and IBM Directory Integrator V6.0. Release 2.6 also provides increased scalability and performance, up to a 400% improvement in performance in some areas of the product.

Workplace Managed Client 2.6

IBM Workplace Managed Client delivers a fully integrated server-managed client to the end user's desktop and can provide organizations with a "rich" front-end for their SOA deployment. It provides flexibility and portability of client-side applications, combined with server-side control and the cost savings traditionally associated with Web-based computing. The Workplace Managed Client features innovative technology that was available on a pilot basis to selected customers during 2005. Version 2.6 is the first generally available release of this technology. It includes support for the ODF standard, which will give organizations a consistent, standards-based format for creating, storing and sharing documents. In addition, the Workplace Managed Client supports a Lotus Notes plug-in for Linux desktops, responding to frequent requests from Lotus customers to provide them with a lower-cost Linux-based desktop option.

Workplace Forms 2.6

IBM Workplace Forms is new electronic forms (or "e-forms") software that can help companies streamline and customize the way business data is captured, processed and exchanged. More efficient than paper forms, IBM Workplace Forms allows customers to integrate back-end corporate data and applications -- such as inventory figures, customer data or pricing information -- with specific formats that provide a unique, electronic replica of paper-based forms, which is critical in vertical industries like insurance, government, banking and healthcare. The new version of Workplace Forms offers support for growing XForms standard and expanded operating platform and language support (internationalization) for worldwide use.

Workplace Designer 2.6

IBM Workplace Designer is a software development tool that allows developers to quickly and easily create new components and applications that can be used with multiple Workplace software products. IBM Workplace Designer makes it easier for developers familiar with Domino Designer -- as well as those familiar with script-based programming languages and Visual Basic -- to build J2EE applications and components. These new tools expand the capabilities of the Lotus Domino developer community, estimated at nearly one million, and help them leverage their skills and investments for standards-based software development.

Using standards such as Java, XML and JavaScript, developers can rely on their existing skills to build components such as expense reports, discussion forums, blogs, and team-based work components. Workplace Designer 2.6 is now integrated with the installation and administration of Workplace Collaboration Services 2.6.

David Kolakowski of Kolaco, a business technology consulting firm, uses Workplace Designer and Workplace Services Express to create custom applications for its customers. "With the introduction of IBM Workplace Designer, Kolaco saw this as an opportunity to take its meeting management system, TrackMeet, and incorporate it into the team collaboration features of Workplace Services Express," he said. "By completely rebuilding the TrackMeet product for Workplace Services Express and using the Workplace Designer API Toolkit, TrackMeet for Workplace adds functionality and integrates with other Workplace software components, including the team calendar and document management."

Pricing and Availability

IBM Workplace Collaboration Services 2.6, IBM Workplace Managed Client 2.6, IBM Workplace Designer 2.6 are available now. IBM Workplace Collaboration Services server sells for $90,000 per processor. IBM Workplace Managed Client is $69 per user. IBM Workplace Forms 2.6 will ship in the second quarter of 2006: The IBM Workplace Forms Server will be $25,000 SRP, IBM Workplace Forms Viewer will be $188 SRP, IBM Workplace Forms Designer will be $649 SRP.

eclipse developer's journal

A DIGITAL PUBLICATION

Your One-Stop Guide to the Eclipse Ecosystem

The main purpose of *Eclipse Developer's Journal* (*EDJ*) is to educate and inform users of Eclipse and developers building plug-ins or Rich Client Platform (RCP) applications. With the help of great columnists, news you can use, and technical articles, *EDJ* is a great information source that you will be able to use to educate yourself on just about anything Eclipse related

Subscribe Today!

FOR YOUR DIGITAL PUBLICATION

(AVAILABLE IN DIGITAL FORMAT ONLY)

6 Issues for $19.99

Visit our site at www.eclipse.sys-con.com or call 1-888-303-5282 and subscribe today!

OFFER SUBJECT TO CHANGE WITHOUT NOTICE

http://eclipse.sys-con.com

SYS-CON.COM

SYS-CON Media, the world's leading publisher of *i*-technology magazines for developers, software architects, and e-commerce professionals, brings you the most comprehensive coverage of Eclipse

BY NOW THERE ISN'T A SOFTWARE DEVELOPER ON EARTH WHO ISN'T AWARE OF THE COLLECTION OF PROGRAMMING TECHNOLOGIES KNOWN AS AJAX!

REAL-WORLD AJAX SEMINAR

DoubleTree Hotel
April 24, 2006
San Jose, CA

The Roosevelt Hotel
June 5-6, 2006
New York, NY

How, in concrete terms, can you take advantage in your own projects of this newly popular way of delivering online content to users without reloading an entire page?

How soon can you be monetizing AJAX?

This "Real-World AJAX" one-day seminar aims to answer these exact questions...

Led by "The Father of AJAX" himself, the charismatic Jesse James Garrett of Adaptive Path, "Real-World AJAX" has one overriding organizing principle: its aim is to make sure that delegates fortunate enough to secure themselves a place – before all seats are sold out – will leave the seminar with a sufficient grasp of Asynchronous JavaScript and XML to enable them to build their first AJAX application on their own when they get back to their offices.

HURRY! LIMITED SEATING THIS SEMINAR WILL SELL-OUT CALL 201-802-3022 TO REGISTER!

Jeremy Geelan
Conference Chair, Real-World AJAX
jeremy@sys-con.com

Register at...
www.AJAXSeminar.com

Early Bird Price (San Jose, CA)	**$695**
(Before March 31 / After March 31 – Price: $995)	
Early Bird Price (New York, NY*)	**$1,295**
(Before March 31, 2006)	
Discounted Price (New York, NY*)	**$1,495**
(After March 31, 2006)	
On-Site Price (New York, NY*)	**$1,695**
(On June 5-6, 2006)	

*Includes SOA Web Services Edge plus Enterprise Open Source Conference & Expo

LIVE SIMULCAST!
AROUND THE WORLD ON SYS-CON.TV

CORPORATE SPONSORS

Adobe telerik BACKBASE TIBCO Microsoft

What you'll walk away with...

A New understading of AJAX and how to make it work for you, plus:

A) **Conference Proceedings Binder**
Full-conference proceedings in a self-contained commemorative binder.

B) ***Ajax in Action* Book**
A copy of the Best-Selling AJAX book by Dave Crane.

C) ***Real-World AJAX "Secrets of the Masters"* Book + DVD**
New AJAX book edited by Dion Hinchcliffe (Release Date: Summer 2006)

D) **Notebook**
Executive style keeps you looking sharp for note-taking wherever your business might take you!

E) **Computer Back-Pack**
A hip new way to conceal your laptop! The Urban Wonder Compu-Pack is an ingenious creation - a 600-denier polycanvas and nylon knapsack!

F) **T-Shirt**
100% preshrunk 6.1 oz. heavyweight cotton tagless shirt features shoulder-to-shoulder tape, and double-needle sleeves and bottom.

G) **Porcelain Mug**
Lightweight and durable, you can use it as a teacup, coffee mug or keep it as a pencil holder.

H) **Satin Silver Contemporary Pen**
Satin-silver plastic barrel with the look and feel of brass. Rubber comfort grip improves writing control and comfort.

I) **Real-World AJAX Seminar on DVD**
Watch complete seminar video and slide presentations (Release Date: Summer 2006)

Featuring... Real-World AJAX Rock Stars! FACULTY

Jesse James Garrett

Father of "AJAX" Who Coined the Term in 2005

Jesse James Garrett is the Director of User Experience Strategy and a founding partner of Adaptive Path, the world's premier user experience consulting company. He is author of The Elements of User Experience (New Riders), and is recognized as a pioneer in the field of information architecture. Jesse's clients include AT&T, Intel, Crayola, Hewlett-Packard, Motorola, and National Public Radio. Since starting in the Internet industry in 1995, Jesse has had a hands-on role in almost every aspect of Web development, from interface design and programming to content development and high-level strategy. Today, information architects around the world depend on the tools and concepts he has developed, including the widely acclaimed "Elements of User Experience" model. He is co-founder of the Information Architecture Institute, the only professional organization dedicated to information architecture. He is also a frequent speaker and writer whose work has appeared in numerous publications, including New Architect, Digital Web, and Boxes and Arrows.

Scott Dietzen

One of the Fathers of WebLogic and J2EE
Ph.D., President and CTO, Zimbra

Scott is widely credited with helping put together the J2EE standard, launching the Web application server category, launching the Java Community Process, and driving the Web services collaboration with Microsoft and IBM. Prior to Zimbra, Scott was CTO of BEA Systems where he was the principal architect of the technology strategy for the WebLogic product family.

Bill Scott

AJAX Evangelist of Yahoo!

Bill Scott, one of the top AJAX experts in the country, is an Interaction Designer and AJAX Evangelist at Yahoo! He is part of the newly formed Design and Practices Team working with teams throughout Yahoo! to create a rich experience on the web. Before joining Yahoo! Bill founded the User Experience Team at Sabre Airline Solutions, part of Sabre Holdings. During that stint he also co-founded Rico, an open source JavaScript framework for creating AJAX & DHTML web applications. Over the past 20 years Bill has been involved in designing and creating user interfaces for video games, military war games, 3D graphics, oil and gas research, software development environments, supply chain planning, and various other scientific and business domains. He posts his musings about user experience on his blog.

David Heinemeier Hansson

Creator of Ruby on Rails

U.S.-based since November, David Heinemeier Hansson is the development lead of Rails, also known as Ruby on Rails, which the official Ruby on Rails website irreverently describes as "a full-stack, open-source web framework in Ruby for writing real-world applications with joy and less code than most frameworks spend doing XML sit-ups." He has had help from a lot of contributors. David is the creator of applications like Instiki, Basecamp, and Ta-da, and has now joined the Chicago-based team of 37signals.com, a privately held company founded in 1999 and committed to building the best web-based software products possible with the least number of features necessary. David will give his first presentation on "AJAX in Rails" at the "Real-World AJAX" one-day seminar.

Adam Bosworth

Vice President of Engineering, Google
One of the Fathers of XML & the Creator of MS Access

Adam Bosworth is Vice President of Engineering, Google. He joined Google in 2005 from BEA Systems, where he was Chief Architect & Senior VP of Advanced Development. Prior to joining BEA, Bosworth co-founded Crossgain, a software development firm acquired by BEA. Known as one of the pioneers of XML, he previously held various senior management positions at Microsoft, including General Manager of the WebData group, a team focused on defining and driving XML strategy. While at Microsoft he was also responsible for designing and delivering the Microsoft Access PC Database product and assembling and driving the team that developed the HTML engine of Internet Explorer 4.0.

Rob Gonda

Bestselling AJAX Author, CTO, iChameleon Group
Editor-in-Chief, AJAX Developer's Journal

Rob Gonda, newly appointed to the helm of SYS-CON Media's AJAX Developer's Journal, is the CTO for iChameleon Group. He is an Advanced Certified Coldfusion Developer, holds a BS in computer science and engineering and an MBA with a specialization in entrepreneurship. He recently wrote a two-part feature article on AJAX for ColdFusion Developer's Journal that in one month became the most-read article in the magazine's history.

Dion Hinchcliffe

Cofounder & CTO, Sphere of Influence Inc.
Editor-in-Chief, Web 2.0 Journal

Dion Hinchcliffe, newly appointed editor-in-chief of SYS-CON's pioneering Web 2.0 Journal, is cofounder and chief technology officer for the enterprise architecture firm Sphere of Influence Inc., in McLean, Virginia. A veteran of software development, Dion works with leading-edge technologies to accelerate project schedules and raise the bar for software quality. He is highly experienced with enterprise technologies and he designs, consults, and writes prolifically. Dion actively consults with enterprise IT clients in the federal government and Fortune 1000. He is a frequent speaker on AJAX, Web 2.0 and SOA and is currently the top-read SYS-CON.com blogger.

Paul Rademacher

Google, Creator of HousingMaps.com

Paul Rademacher is the creator of HousingMaps.com, which combined Craigslist and Google Maps for the first web mashup. Paul holds a Ph.D. in Computer Science from UNC-Chapel Hill, and worked as an R&D Engineer at Dreamworks Animation on such movies as Shrek 2 and Madagascar. Since creating HousingMaps, Paul is now at Google.

Ross Dargahi

Well-known AJAX Evangelist
Co-founder and VP of Engineering, Zimbra

Ross Dargahi, a co-founder of Zimbra, is presently the company's VP of Engineering. Prior to Zimbra, he was Director of Engineering and Director of Product Management with the Messaging Products Group at Openwave Systems - where he built and led the engineering team that designed and architected large-scale messaging subsystems and was responsible for unified messaging, multimedia messaging (MMS), and voicemail. Ross joined Openwave as part of the 1999 acquisition of Onebox where he was a founding engineer and senior architect. Prior to Onebox, Ross was at Sun Microsystems' JavaSoft division where he was a lead engineer with the Java Server Group responsible for network computer and embedded server products.

Dave Crane

Co-author of Bestselling AJAX Book 'AJAX in Action'
Dave will autograph his bestselling book for all delegates!

Dave Crane is senior developer and architect for Historic Futures Ltd., a UK firm specializing in web-based supply chain management solutions. He is also lead author of the bestselling 'AJAX in Action', and has been using the technologies that we now know as AJAX in production for several years. He has worked extensively in the IT industry over the last ten years, in areas as diverse as home automation, finance and national policy planning, for companies ranging from start-ups to FTSE 100 market leaders, using technologies as diverse as J2EE, Python, Ruby, PHP, Linux and .NET. He holds degrees in environmental science, parallel computing, and a Ph.D. in simulation modelling and philosophy. He is active as a writer, trainer and mentor in the AJAX commun

Christian Cantrell

Coauthor of JavaScript/Flash Integration Kit
AJAX/Flash Integration Guru

Christian Cantrell is a Product Manager for Developer Relations with Adobe and a big fan of AJAX. He has been developing large-scale, web-based applications in ColdFusion, Java, JSP, and Flash for more than six years. He is the author of numerous tutorials and white papers, and is coauthor of Flash Enabled: Flash Design & Development for Devices as well as the JavaScript/Flash Integration Kit. Christian is a leading expert on AJAX to Flash integration.

Sahil Malik

telerik Tech Evangelist, Microsoft MVP
Author of Bestselling ADO.NET 2.0 book

Sahil Malik is an independent consultant, trainer and mentor in various Microsoft Technologies and has been closely involved with telerik as a technology strategist. He has worked for many large clients across the globe including a good deal of Fortune 100 companies and US government organizations. He is currently leading the office of Emerging Technologies at a prominent government office where he is in charge of reviewing, assessing and recommending various technologies to support the organization. Malik frequently speaks on a variety of .NET related topics at local user group meetings and industry events. For his community involvement and contribution, he has been awarded the Microsoft MVP award. He can be reached at www.winsmarts.com.

Jouk Pleiter

Co-Founder & CEO of Backbase

Jouk Pleiter is the CEO of Backbase, a leader in the field of Rich Internet Applications and AJAX development software. Backbase's clients include ING, ABN AMRO, TNT, KPN, Comsys and Heineken. Backbase operates globally with offices in San Mateo (North America) and Amsterdam (Europe). Since 1995, Jouk has been an entrepreneur: he founded three successful software companies. Prior to Backbase, Jouk was part of the founding team at the web content management company Tridion, where he led the product management operations, and was driving the company's efforts to become a leader in the European WCM software market. Jouk previously was part of the founding team at the Interactive Agency Twinspark where he grew the company to a leading market position in Europe and was instrumental in the sale of Twinspark to Agency.com. He has an MBA from the University of Groningen.

Kevin Hakman

Director of Product Marketing for TIBCO
General Interface TIBCO Software

Kevin Hakman is the director of product marketing for TIBCO General Interface, the award winning AJAX and Rich Internet Application framework and toolkit. Kevin Hakman pioneered AJAX in the enterprise co-founding General Interface in 2001. Since that time General Interface (aka 'GI') has been powering Web applications that look, feel and perform like desktop applications, but run in the browser at Fortune 500 and U.S. Government organizations. General Interface was also the first to use its own toolkit to provide full visual tooling for AJAX when it released it's 2.0 Version in 2003. TIBCO acquired General Interface in 2004 to extend its vision for service oriented applications to the end user. Kevin is a contributor to the SOA Web Services Journal and the AJAX Developer's Journal.

Shanku Niyogi

Product Unit Manager of the UI Framework and Services Team Microsoft Corporation

Shanku is Product Unit Manager of the UI Framework and Services (UiFX) team, which is responsible for delivering high-productivity UI framework technologies for the .NET platform, including ASPNET, Atlas, Windows Forms, and frameworks for smart clients. Prior to his current role, Shanku was Group Program Manager of the Web Platform and Tools team on the Whidbey release of ASPNET and Visual Web Developer. Shanku joined Microsoft in 1998 as a developer, having spent several years shipping products in the Windows ISV industry. Shanku holds a Bachelor of Mathematics degree in Computer Science from the University of Waterloo.

For more information...

Call 201-802-3022 or email events@sys-con.com

SYS-CON EVENTS **For more great events visit www.EVENTS.SYS-CON.com**

NOTE: SPEAKER LINE-UP SUBJECT TO CHANGE WITHOUT NOTICE
VISIT WWW.AJAXSEMINAR.COM FOR THE MOST COMPLETE UP-TO-DATE INFORMATION

Open AJAX Is Loosed Upon the World

"Pretty. But Can They Fight?"

BY ROGER STRUKHOFF

IBM's announcement last month of the Open AJAX initiative brought to the fore a few key points. The two most obvious are that IBM remains an 800-pound gorilla that the other primates will follow through the business jungle, and that AJAX has emerged as an ironically disruptive technology.

Ironic in that AJAX comprises not only a seemingly modest collection of modest elements – XML and JavaScript – but that these elements are hardly new, work in a loosely coupled fashion, yet represent the leading edge of what's going on with application development in the world today.

Open AJAX also brought in the heavyweights Oracle, Novell, Red Hat, Google and Yahoo, in addition to solid middleweight Borland and major featured performers Zend, Zimbra, Laszlo, Mozilla, and the Eclipse Foundation. Prominent in their absence were Sun, HP, and Tibco, although spokespeople from both Sun and Tibco have told me that they haven't ruled out future participation in this coalition.

Like the applications upon which it is focused, Open AJAX is loosely coupled, with no formal organization or management structure. It represents, in the word of one member executive who spoke to us, "IBM putting its foot down and the rest of us destined to follow its footsteps."

The essence of AJAX is its ability to keep applications from returning to a central server every time a user wants to know something. The theory is that this not only makes things easier for the server and the network, but also enhances the user experience. Things are speeded up, server and network loads reduced, and happiness ensues all around.

But by employing the venerable XML and the increasingly venerable JavaScript, AJAX also represents the continued movement away from proprietary software (including application development environments) that has been witnessed over the past 18 months. Some analysts argue that IBM's leadership and support of Open AJAX is simply a defensive move to enable it to take a measure of control over the many wild open-source horses that have left the barn.

IBM's lead on this project, David Boloker, doesn't think so—you can learn more about his views by visiting www.sys-con.tv and finding a recent interview we did with him as Open AJAX was being announced. And many executives within this initiative supported Boloker's leadership, many of them calling it "long-term" and "tremendous."

This effort seems to be not so much a defensive marketing move from IBM as a provocative leadership move driven by technologists rather than marketers. Many of the participants add a particular flavor (or "personality" in Boloker's term) to this overall initiative, personalities that will in fact be competing with one another for developers' favor.

As these personalities compete for one another within a defined AJAX universe, developers may get the best of both worlds: the latest and coolest way to develop and deploy apps in an environment that is neither anarchic nor held hostage by a single vendor.

"Pretty, but can they fight?" as Donald Sutherland's character asked a generation ago in Kelly's Heroes. (If you get this allusion, then thanks for putting on your reading glasses to take in this column. If you don't get it, Donald looks remarkably like his son Kiefer but plays goofier characters.) Whether Open AJAX will in fact be able to fight its way to AJAX leadership in coming months, and whether this will be a net positive for IBM will be, as they say, decided by the market.

The real test will be to see whether JavaScript-driven local performance becomes a central application development issue, or just a nice feature that will be integrated here and there. Google maps is cool and all that, but not always accurate, and hardly made of the stuff that heavy financial, government, industrial, or any other mission-critical apps within any other vertical market demand.

But the initiative is not static, and its lack of formality will presumably let all the players adapt as they see fit, and thereby not lock in developers to any particular flavor or personality, all the awhile affording many risk-averse companies and agencies a measure of security. After all, no one gets fired to this day for buying IBM, so only time will tell whether no one will get fired for embracing Open AJAX.

ABOUT THE AUTHOR

Roger Strukhoff, editor-in-chief of *WebSphere Journal*, is West Coast Bureau Chief for the SYS-CON News Desk, and President of www.wdva.com. He spent 15 years with Miller Freeman Publications and The International Data Group (IDG), then co-founded CoverOne Media, a custom publishing agency that he sold in 2004. His work has won awards from the American Business Media, Western Press Association, Illinois Press Association, and the Magazine Publishers' Association. Read his blog at http://www.rssblog.linuxworld.com.
roger@sys-con.com

The Developer Paradox:

No time to test your code?
But long hours reworking it & resolving errors?

Check out Parasoft Jtest® 7.0
Automates Java testing and code analysis.
Lets you get your time back and deliver quality code with less effort.

- **Automated:**
  Automatically analyzes your code, applying over 500+ industry standard Java coding best practices that identify code constructs affecting performance, reliability and security.

  Automatically generates and executes JUnit test cases, exposing unexpected exceptions, boundary condition errors and memory leaks and evaluating your code's behavior.

  Groundbreaking test case "sniffer" automatically generates functional unit test cases by monitoring a running application and creating a full suite of test cases that serve as a "functional snapshot" against which new code changes can be tested.

- **Extendable:**
  Industry standard JUnit test case output make test cases portable, extendable and reusable.

  Graphical test case and object editors allow test cases to be easily extended to increase coverage or create custom test cases for verification of specific functionality.

- **Integrated:**
  Integrates seamlessly into development IDE's to support "test as you go" development, and ties into source control and build processes for full team development support.

To learn more about Parasoft Jtest or try it out, go to www.parasoft.com/JDJmagazine

PARASOFT®
We make software work.™

Automated Software Error Prevention™

Java 911

Put out code fires with JProbe®

Tired of constantly fighting fires? Extinguish code issues and prevent future memory flare-ups with Quest Software's JProbe.

Discover the root cause of performance and memory issues in your Java applications. Easily test and fix your applications without any code changes. And even pinpoint performance problems down to the line of code. All with award-winning JProbe.

Stop the fire drills. Gain confidence in your code with JProbe.

Download a trial of JProbe at:
www.quest.com/911