SYS-CON

DEPLOYMENT • TELECOMMUNICATIONS • GRID

WebSphere® JOURNAL

www.SYS-CON.com/WebSphere

The World's Leading Independent WebSphere Resource

JANUARY 2005 VOLUME 4 ISSUE 1

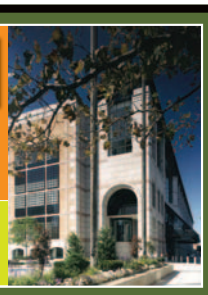

DISPLAY UNTIL MARCH 31, 2005

$8.99US $9.99CAN

0 09281 03422 3 01>

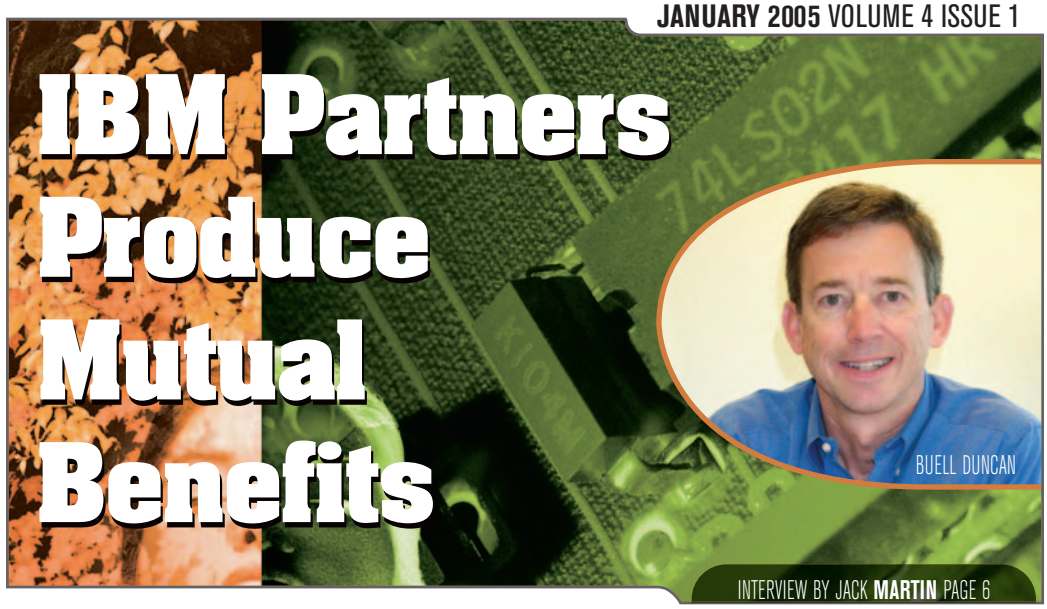


Middleware is Everywhere.
Can you see it?
IBM
3
4
1
5
2
Tivoli
Key
1. Star footballer heads it in at buzzer.
2. Fans call rival fans to gloat.
3. 300% more calls flood network.
4. IT systems sense and respond instantly.
5. Telecom scores big with customers.
MIDDLEWARE IS IBM SOFTWARE. It's software that automates your IT systems to align with your unique business goals. Open, flexible, end-to-end solutions that help you keep up with customer demands on your infrastructure. It's how to avoid costly interruptions in service and fulfill your service level agreements. It's how to stay on top of your game, no matter what size you are.
Middleware for the on demand world. Learn more at ibm.com/middleware/agility
ON DEMAND BUSINESS
IBM, the IBM logo, Tivoli and the On Demand logo are registered trademarks or trademarks of International Business Machines Corporation in the United States and/or other countries. © 2004 IBM Corporation. All rights reserved.

Will Every Child Be Left Behind?

BY JACK MARTIN

The politicians are at it again with their general collective hallucinations. Recently Rod Paige, the U.S. Secretary of Education, made the following statement:

There is a new fervor in American education and a new creativity that's being driven in part by this generation of tech-savvy students. We are already seeing some remarkable results, and I believe this trend bodes well for the future of our country. As the report noted, America's students are our ultimate constituents, and we need to listen to them.

Paige added that teachers are transforming what can be done in schools by using technology to access primary sources, exposing students to a variety of perspectives, and enhancing students' overall learning experience through multimedia, simulations, and interactive software. At the same time, teachers, principals, and administrators are able to better track student achievement and adjust instruction more effectively to individual needs.

The report includes Paige's vision and recommendations for a National Education Technology Plan based on input received from educators and technology experts across the country.

According to the report, the technology that has so dramatically changed the world outside our schools is now changing the learning and teaching environment within them. This change is driven by an increasingly competitive global economy and the students themselves, who are "born and comfortable in the age of the Internet."

"As these encouraging trends develop and expand over the next decade, facilitated and supported by our ongoing investment in educational technology..." the report said, "...we may be well on our way to a new golden age in American education."

Well Mr. Paige, if you took the time to visit some real public schools and spoke with some real students and teachers, you would find that the average public school teacher in America has almost no understanding of technology nor an ability to use it for any useful purpose except in the most rudimentary fashion. School teachers as a group have been left behind in the technology revolution just as the weavers were in 19th-century Europe.

If you do not believe me, ask a teacher to explain what an application server is. Or ask something much simpler, such as, how does an operating system work? You'll be greeted with a blank stare. Or try to get a group of them to explain how a computer video game works and, more importantly, why almost every one of their students find those games more interesting than anything the public school has to offer.

The reality is that the average public school teacher doesn't know anything about XML, Web services, Grid computing, application servers, or almost any important technology, no less have the ability to explain and teach it to someone.

I have asked many public school teachers why they do not have a basic understanding of information technology. (By the way, whenever I have used the term information technology, I had to explain IT to them, and they always responded, "Oh, computers!") And the conversation always follows with the statement that they were never trained "on computers." The conversation then turns to how bright the kids are with computers – they use the Web, chat, send e-mail, and play computer games all the time.

I have tried and failed to explain to many teachers over the years that using an application (if the application was created with care and intelligence) requires very little intelligence and few skills, but it does require patience. It is *writing* an application that requires high skills and intelligence. This is where the conversation ends and they change the topic or walk away.

Mr. Paige, information technology skills come to those who go out and get them, not to those who wait for someone to hand deliver them. In order for anything you have said to ever be true, you first have to re-educate all of the public school teachers, and then take it from there.

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

WebSphere® JOURNAL

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

EDITOR-IN-CHIEF JACK MARTIN
EXECUTIVE EDITOR GAIL SCHULTZ
CONTRIBUTING EDITOR PATTI MARTIN
PRODUCT REVIEW EDITOR JAY JOHNSON
EDITOR NANCY VALENTINE
ASSOCIATE EDITOR JAMIE MATUSOW
ASSISTANT EDITORS NATALIE CHARTERS, SETA PAPAZIAN

WRITERS IN THIS ISSUE
Moon J. Kim, Max King, Colm Malone, Jack Martin, Ellen Matheson McKay, Dikran S. Meliksetian, Mike Perera, Barry Searle, Joe Winchester

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 STEVE MICHELIN
ACCOUNTS PAYABLE BETTY WHITE
ADVERTISING DIRECTOR ROBYN FORMA
ADVERTISING MARKETING DIRECTOR DENNIS LEAVEY
ADVERTSING SALES MANAGER MEGAN MUSSA
ADVERTISING SALES MANAGERS KRISTIN KUHNLE, DOROTHY GIL, KIM HUGHES
LEAD DESIGNER LOUIS F. CUFFARI
ART DIRECTOR ALEX BOTERO
ASSOCIATE ART DIRECTORS RICHARD SILVERBERG, TAMI BEATTY, ANDREA BODEN
WEB DESIGNERS STEPHEN KILMURRAY, MATTHEW POLLOTTA
ONLINE EDITOR MARTIN WEZDECKI
CIRCULATION SERVICE COORDINATORS EDNA EARLE RUSSELL, LINDA LIPTON, MONIQUE FLOYD
CUSTOMER RELATIONS MANAGER BRUNI STAROPOLI

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

WORLDWIDE NEWSSTAND DISTRIBUTION
CURTIS CIRCULATION COMPANY, NEW MILFORD, NJ

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

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

Why is it that some Java guys are more relaxed than others?

These days, with everything from customer service to sales and distribution running on Java, keeping your enterprise applications available can be pretty stressful.

Unless of course, you've discovered the power of Wily. No other software offers our level of insight. Which means you'll be able to monitor transactions and collaborate with colleagues in real-time. Even more important, you'll be able to optimize performance across the entire application—end-to-end.

So put Wily to work. And keep performance problems from pushing you over the edge.

Get Wily.™

1 888 GET WILY | wilytech.com

wily technology

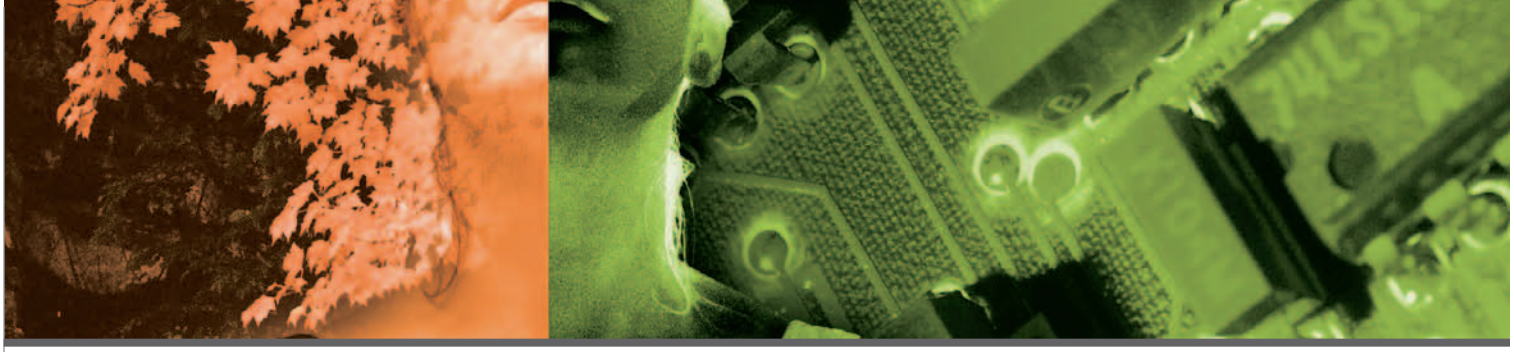

IBM Partners Produce Mutual Benefits

BUELL DUNCAN, GENERAL MANAGER OF IBM ISV & DEVELOPER RELATIONS

The elements that compose a successful relationship

INTERVIEW BY JACK MARTIN

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

Jack Martin, editor-in-chief of *WebSphere Journal*, recently sat down with Buell Duncan, general manager of ISV and Developer Relations for IBM.

WJ: How do you spend a typical day?
BD: That's a very good question. Over the course of any given week, I spend a third of my time working with the IBM team developing initiatives and programs and considering investments to strengthen our relationship with the software community at large. I spend probably the majority of my time each week, over 50% certainly, working with software companies, with developers, and with the venture capital community.

WJ: When you say you work with the venture capital community, do you mean the outside VCs?
BD: The outside VCs. We do spend a lot of time developing programs and industry-specific offerings across the world industry network. We are working with the different brands on strategic alliances with certain software firms, along with the SMB organizations for the ISV Advantage Program, which is very successful, as we reach local and regional software companies. When you talk about all the time I spend internally working across IBM, it is very important, Jack, to emphasize that although I report to Steve Mills, the mission of this team is global and it is across the company, meaning representing the systems group, software group, global services, and of course, the SMB organization.

WJ: That's an enormous job.
BD: I wouldn't say that. It's a fun job, it's an exciting job, and what's most exciting about this job is the breadth, as well as the amount of time that you have to spend and that you want to spend understanding the requirements and the opportunities of the software firms themselves. I met with the CEO of Relavis, an IBM Business Partner yesterday. This morning I was on the phone with the head of the software business at CSC, Computer Sciences Corporation. Every day this job gives you the opportunity to work directly with outside companies.

WHAT IS AN ISV?
An ISV (independent software vendor) is a software company that creates products that run on another company's platform.

WJ: When you say you work with them, what does that look like?
BD: Let me paint you the broader picture, if I can. As you know, we made a decision as a company five years ago to exit the application business.

WJ: Absolutely.
BD: We made that decision because it was not in accord with IBM's overall strategy, and because frankly we had been at it in one form or another for a number of years. We believed that by building strong relationships with application providers we could be more successful than if we tried to pick a spot and do it on our own. IBM's strengths are around client relationships, around technology, and around our services capabilities. As the marketplace is increasingly looking for solutions, a key component of that solution is always going to be an application asset, so we started five years ago with the first strategic alliance where we made specific commitments to Siebel. They in turn made specific commitments to IBM in terms of technology enablement. For example, WebSphere is obviously a key element in terms of go-to-market coordination and activities together.

WJ: Very interesting.

BD: Since that first strategic alliance five years ago, we have now done 2,500 customer engagements with Siebel. That's pretty significant, and we have seen our share of business associated with Siebel license revenue grow consistently year in and year out, versus our competitors.

We started with a series of strategic alliances with global leaders like SAP, Siebel, and i2, and then we expanded that into what is now approximately 100 strategic alliances with large software firms around the world. These companies include i-Flex, building technology solutions for the financial industry, and SAS, with whom we have a relationship in the health care/ life sciences space. However, it wasn't enough to have 100 strategic alliances because there are literally thousands of software firms that customers are interested in, and the solutions they bring to market.

Last year we introduced a program that works hand in hand with the SMB organization led by Marc Lautenbach. We launched the ISV Advantage Program where we focused on 300 of the top local regional ISVs around the world. These ISVs are leaders in the application space in their geographies such as Munich, Milan, Singapore, or Minneapolis. We've made commitments to help these ISVs technically migrate to IBM platforms – WebSphere and DB2, and we've seen our share of their business grow dramatically. In the past 18 months, 250 companies have joined this program and have signed agreements with us to lead with IBM in the marketplace, delivering the majority of their business on IBM platforms. Among these 250 companies in the ISV Advantage Program for SMB, we've seen their enablement on WebSphere grow. In 2003, 24% of these companies were enabled on WebSphere. In 2004, that number tripled to 78% of companies who have their applications fully enabled on WebSphere, and 96% of them, as we speak, are committed to doing that. We still have a few to go, but it shows you the kind of connection points that we are making with these companies. The theory is pretty straightforward that these application companies lead with IBM , if we work together in the marketplace. The more successful they are, the more successful we'll be.

That would suggest [if you added up the strategic alliances we have with approximately 100 global players, and in many cases those strategic alliances have an IGS or BCS (Business Consulting Services) practice assigned around them] that we have contracts with about 500 software companies around the world that are our partners of choice and we are their partners of choice.

WJ: Wow!

BD: That's still not enough. There are thousands of companies, and every day we are finding there are new entrants in the marketplace. So, how can we reach a much broader number of software firms? In 2004, we launched the Partner World Industry Network Program. We will add five more industries this year (2005). This will be the flagship program that we will drive for the coming years. In the first eight months that this program has been in the market, more than 2,200 software companies have joined. When they join they have access to specific industry content and resources. In other words, the way we go to market with our clients in banking, health care, or any other vertical industry is that we provide the technical blueprints and roadmaps for enablement, and give them resources to help enable their applications to our technology like WebSphere, DB2, Lotus Workplace, and orchestration and provisioning from Tivoli.

One of the most important things we do, besides provide these companies with technical support and enablement capabilities, is offer them marketing benefits when they have completed the enablement and when they have their first customer implementation. These marketing capabilities include deep discounts on advertising. We have agreements with 220 publications around the world. Members of IBM Partner World Industry Networks who are enabled and who have customer references can get discounts in the 70-80% range on advertising their solutions with IBM in these publications. We have a demand-generation campaign and collateral materials that we make available to these companies.

Most important, what we do, in terms of the feedback we get, is provide the ability to connect these very small firms to the IBM sales force within our Business Partner organizations. So, when they have a lead or an opportunity, they are enabled on WebSphere and pSeries, and want to find out who to work with to close this and get help from IBM, we can connect them directly with one of our sales people or one of our business partners to close the transaction. We have seen some 200 sales connections take place already in just the last 60 days since that offer of programs was made.

IBM's channel strategy consists of a spectrum holistic approach, working with the global players like SAP and Siebel, in terms of getting more and more share of their license revenue on IBM platforms. We've been able to bring the technical and enablement support and go-to-market benefits to literally thousands of companies who want to partner with IBM.

Working with the venture capital community, we recently completed a tour of six cities around the world including London, Sao Paulo, Washington DC, Shanghai, Palo Alto, and Boston where we held roundtable sessions. For example, in Palo Alto, some 50 venture capital firms attended the session and each of them brought portfolio companies. We talked about ways in which we could get connected to build up solutions. We need to build out solutions for our customers, which involves working with ISVs and application providers. How do we get better connected with these ISVs, so we can go to market together and provide stronger solutions to the mutual end-user customer? As we do this, everyone wins. IBM is in a stronger position because of the breadth and strength of our solution portfolio, its access to market for the smaller startup software companies and of course, it benefits the VCs who are funding these companies.

Buell Duncan is general manager of IBM ISV & Developer Relations. He is responsible for IBM's alliances and programs with independent software vendors and solution developers. Mr. Duncan, a graduate of Vanderbilt University in Business and Economics, joined IBM as a marketing representative in 1975. He held various management positions, including general manager, iSeries; vice president, Sales, AS/400; general manager, International Operations, Worldwide Sales and Services; general manager, Systems Sales, EMEA; and general manager, Industries, IBM North America.

The real beneficiary though is the end-user customer because of the breadth of the solutions that they now have available to them in dealing with IBM.

WJ: How do you choose a VC? How do you decide if you are going to work with a specific VC?
BD: We have a dedicated team at IBM, Venture Relations, whose mission in life is not to go out and fund companies, but to work with the leading venture firms in the industry. We are interested in two things and this is very similar to our approach to working with ISVs. We are interested in those firms who are leaders in their markets such as Kleiner Perkins or Hummer Windblad Venture Partners. Ann Windblad was with us in Palo Alto for the recently held VC session. We're very interested in building stronger relationships with the leaders in the industry within their particular markets, and with those who have access to innovative technology firms that we believe will be complementary to IBM's effort. We are also interested in those companies that want to work closely with IBM, that believe as we do in the importance of open standards and in the value of a true partnership.

Several things set IBM apart from other companies. First, the way the market is evolving seems to be playing very much to our strengths. The primary fact that sets IBM apart is the total commitment to open standards.

Commitment to open standards. Partnering versus competing. Investing in go-to-market support for our partners.

If you're an application company and you're writing an application, why would you ever write an application that runs only in a Windows environment when you could develop that application in a Java environment and it could be deployed in a Windows, Linux, UNIX, or legacy environment? IBM's complete commitment to, and support of, open standards (that means Java, Linux, and Eclipse) is something that is gaining real traction in the marketplace with our partners, including the VC community.

The second thing that sets IBM apart is that we are genuinely committed to partnering. We are not in the application business. We are not confused about whether we are competing or partnering. We built a strategy completely around partnering with not just one or two companies, but literally thousands of software firms.

WJ: I think it's undeniable that you have the biggest business partner program in the world.
BD: What sets it apart is the fundamental belief in open standards, that's number one. Number two is the belief in partnering versus competing. Third, the feedback we are getting very consistently from the analyst community is that what we are doing with these companies goes beyond just providing technical support on the go-to-market side; it's unique. In other words, this idea of helping smaller firms in terms of sales connections, closing deals together in the territory, and helping them generate business sets IBM apart. Thus, due to these three items, we are gaining real traction right now in the marketplace.

WJ: What qualities must a company possess to be a candidate for a strategic alliance partnership with IBM?
BD: First of all, we are constantly evaluating the needs of customers and thinking, "what are the solutions that best address those needs?" Solutions that our customers are interested in are a combination of world-class technology, services capabilities, and software applications. So each business – whether it is the banking business in IBM, the eServer pSeries business, the WebSphere brand business, or the SMB markets – is looking at the customers and the solutions that drive business opportunity. As a result, the ISV applications are critical in every case. We work with ISVs that have the leading applications in a given area. The given area may be for risk and compliance in banking or it could be in the area of health care and life sciences–specific applications for data mining. So, who has the leading application in that space, and who is committed to open standards in working with IBM to address those opportunities? Those are the two priorities.

WJ: When you look at the ISV Advantage partners, what do they have to have to be qualified to be there?
BD: Frankly, what I just described to you regarding the leading application providers in their market segment, and who's committed to partnering with IBM and leading with open standards applies across all of our partners.

Also, if you were to ask Marc Lautenbach what companies he's looking for, the only spin he would put on it is their presence in the local market, so that he can connect his sales team locally with those ISVs.

WJ: So, correct me if I'm wrong, it sounds like what IBM has done is leveled the playing field for

ISVs so that whether you are an enormous ISV or a local ISV, you can get all the benefits of working with the largest technology company in the world.
BD: Correct! Obviously, we have larger investments around the larger companies. For example, we have competency centers in place that are physical centers with people staffing them for SAP and Siebel.

We also have BCS consulting practices dedicated to some of the largest ISVs, but the fundamental benefits in the capability to work with IBM are available for all, assuming two things. One, they embrace open standards and two, they have a solution or application asset that is a leader in their space.

WJ: So, what you are describing, correct me if I'm wrong, is an ISV starting out today that started as a PartnerWorld member. If they had an application that they wrote that got traction, they could leverage off the IBM scale. Is that what you are saying – that as they grow, you'll grow?
BD: That is correct.

WJ: Could you describe from inception to somebody who's rockin' and rollin,' how you would see that happening for them?
BD: For a company that we may be introduced to by a venture firm, what we are looking for is those companies that have solutions, not just an idea, but a solution that they have sold to at least one customer, and they want to tighten the relationship with IBM. What is it that they need to do?

WJ: Correct.
BD: What they need to do is join Partner World Industry Networks, an initiative that 2,252 companies have joined in the last eight months. These ISVs will have access to our strategy to work in a given industry. They'll have access to a technical road map of what we are looking for in terms of the solutions enabled to our products that we like to take to market in that industry. We will provide access to our IBM Innovation Centers to help them migrate those applications. We will assign technical architects to work with them, to build their product road maps and then watch as their application is enabled to the appropriate technologies for that given industry, which would include WebSphere, as an example. Then we work together to drive demand for that solution.

How do we do that? Here's an example. There's a company called Vertical-i. Vertical-i is in the health care and life sciences space, and they ran a joint ad with IBM with the business partner logo and a joint ad with us in a publication called *BioIT World*. They were fully enabled on WebSphere and IBM eServer offering. We talked about their solution, got that out in front of tens of thousands of readers, and they were able to run this advertisement as an IBM Business Partner at a cost that was 70% cheaper than if they were trying to do it themselves.

Another example is a small company called RJS Systems, based in Minnesota, enabled on IBM technology. RJS had a business opportunity in upstate New York and contacted Sales Connections. We assigned an IBM salesperson and an IBM Business Partner to work with RJS out of Minnesota to close a business transaction with an insurance company in upstate New York. These benefits are available today to these companies, and as their success builds, they will be identified by Marc Lautenbach's team as a leader in their space in their local market and begin to get some of the benefits and the ability to join the ISV Advantage Program. The benefits include sitting down together and doing marketing workshops, jointly funding demand-generation campaigns, and working to track our success and to close more business together in the market.

We are being told by our partners, customers, and obviously by analysts, that IBM has developed arguably the most comprehensive set of programs and the ability to connect to ISVs of any company in the industry because of the momentum behind open standards and IBM's leadership there, and because of our commitment to the partner and our actually going to market with these companies, whether they be small, midsize, or large.

WJ: Right now, where do you see a shortage of ISVs; where is there demand for something, but there's not enough of them to go around? There are literally tens of thousands of innovative companies coming up with new solutions every day in the application software space. But what are the things you are looking at as you look ahead through 2005


BUELL DUNCAN

and beyond? What is your focus going to be?

BD: I would point out three things that we believe are very important for us as we move ahead. Number one, this is expanding the number of industries, so we'll announce new industries in the first quarter of 2005. In addition to expanding the number of industries we will increase our investments and our focus on the emerging markets' high-growth countries, in particular, China, India, Brazil, and Eastern Europe. There is an enormous opportunity for high growth and innovation in countries like India. Obviously, new software development is occurring. The development isn't just for the local market, it's also development that impacts the rest of the world, so we are going to increase our efforts there even further. The third area is one that I think is particularly exciting in terms of looking ahead, and that is the shift in the way software is delivered from the traditional license revenue model of most traditional ISVs to software as a service or a subscription delivery. I think that IBM can bring a lot to the table to help those companies that are looking at transforming their existing businesses from license models to subscription models, or those companies that are born on the Web and are starting as net native around the subscription delivery models from day one. They are looking for strong partners to help them drive demand for their solutions in the marketplace, to provide hosting capabilities, and again, who fully believe and embrace open standards and open source. We have a real opportunity to work with those companies as we are seeing such growth around software as a service.

WJ: So how do you measure all of this?

BD: We measure it in terms of the share of business that we get working with these companies. For example. we've seen a 5% share gain in the past year on our servers associated with SAP. We've seen a 10% share gain of global services associated with Siebel's business. Those are examples of where we measure share of transactions in the marketplace. We also measure revenue associated with these partnerships, and we've seen very strong growth there. We believe the ISV Advantage Partners themselves, the local regional companies, drive over half a billion dollars of revenue for IBM because of these relationships.

Finally, when you look at the broader market, we measure through some very, very detailed research. We measure the consideration and preference of our ISVs more broadly, as well as who is their preferred partner. In the last 18 months, we have seen double-digit gains in terms of consideration and preference for IBM, while our data suggest at the same time the ISV community in terms of consideration and preferences is declining by a similar amount as it relates to Microsoft. Why is that? Because the market is moving to open standards due to the strength of IBM's product portfolio (again, WebSphere is the market share leader), and because of the set of programs and the commitment we have in partnering with these companies.

WJ: Final question. For the person just starting out, who has his or her application written and either has gotten the first customer or is reading this interview, where does he/she go to find out the first step to starting with you?

BD: On one end of the spectrum are the global players, like SAP and Siebel and the industry leaders like Siemens. Then keep going down the spectrum to the ISV Advantage local regional ISV, and take it to the thousands of other companies who are bringing innovative solutions and technology to market, typically in an industry-by-industry approach. Hopefully what I have conveyed to you in the last half hour is how we cover that. Remember the question that I asked earlier, though? Who in his right mind would ever develop an application that runs only in a Windows environment?

WJ: Sure.

BD: Shouldn't they have thought about that when they developed it in Java and wrote in multiple environments? The only reason a company would ever do that is because they don't have the availability of skills, and so when you play out this spectrum, I really want to emphasize what we've also done over the past several years to build up the relationship with the software development community through our developerWorks initiative. developerWorks lets us reach over four million software developers around the world. In fact, 4.5 million of the estimated 10–12 million software developers in the world have registered on developerWorks for content, skills, tutorials, downloads, and more information about XML Web services, WebSphere, and DB2. Building a tight relationship based on the skills, awards, and recognition of this site, is the industry-leading aspect of developerWorks.

We've also done outreach programs around the world. This year, we've had over 55,000 software developers attend two-day classes in China and India, in the U.S. and Europe, and in Latin America to build up their technical skills around IBM offerings and around open standards. Now, the last stop on this spectrum if you've got this picture in your mind is, it's not enough to work just with the professional developer community from client, customers, and software firms, but we are now reaching out to the universities. Our academic initiative focuses on course curriculum at the university level on a global basis to teach these fundamental skills around open source and around IBM offerings, again using WebSphere and DB2 as an example. When you look at this holistic view from end to end, you can see this is an investment that IBM is fully committed to across the entire corporation, and it's something that we've been asked for some time and we'll stay at it, because we believe in the end it has a very, very significant impact on the success of IBM's overall business.

What I really want to encourage our partners to do is to get connected, because what we are driving at so hard is building out a much broader portfolio of partners than we have today, because we believe that the business opportunities around on demand are there. I encourage them to go to IBM.com/ISV or IBM.com/developerworks.

Combining AEP Practices for Effective Error Prevention

As I mentioned in last month's Buzz, the AEP practices of coding standards and unit testing should be considered two sides of the same coin. Each practice can uncover problems that the other cannot. Consequently, to identify and prevent a wide range of software problems, it's important to perform *both* coding standard analysis and unit testing, not just one or the other.

This month, I'd like to elaborate on how combining specific AEP practices is the most effective way to identify and prevent specific errors. For example, let's consider resource or memory leaks. Application-level runtime error detection and/or profiling can identify leaks, but typically not until late in the development process, when correcting the leak is typically time-consuming, difficult, and expensive. With the appropriate technologies, leaks could be detected at the unit level. However, unless the development team is required to close all resources that are opened within a unit, a leak found at the unit level might not indicate a true application-level leak.

One effective solution to this dilemma is to combine coding standard enforcement, unit-level leak detection, and application-level leak detection/profiling. First, implement a coding standard that requires that the same class or function which allocates a resource must also deallocate that resource. By checking compliance with this coding standard, you can ensure that the team's developers write code in a way that prevents leaks. If this coding standard does not make sense in all situations, individual exceptions can be documented and suppressed. You can then use unit-level leak detection to dynamically verify that leaks do not occur within the unit. After all units are developed in this manner, you can use application-level leak detection and/or profiling to verify whether leaks occur at the application level. When the team works in this manner, significantly fewer leaks reach the application level, where they are more time-consuming, difficult, and expensive to fix.

The Parasoft AEP Methodology is constructed to combine multiple practices in this supportive manner. It encourages development teams to determine which types of errors they need to prevent, then combine practices to create a "leak-proof" system for identifying and preventing those types of errors. For details about this methodology, visit **www.parasoft.com** .

– Adam Kolawa, *Ph.D.*
Chairman/CEO of Parasoft

Part 1 – Basic concepts

Automated Deployment of Enterprise Application Updates

BY: BARRY SEARLE & ELLEN MATHESON McKAY

Barry Searle is the architect for WebSphere Tools for Automated Build and Deployment. A professional engineer, he has worked at the IBM Canada Lab for over 15 years on various application development tools. Prior to that he worked on developing command and control systems, and leading complex communications development projects.
searle@ca.ibm.com

This two-part article discusses application deployment, particularly automated updates, to IBM WebSphere Application Servers in a large-scale enterprise environment. It applies to versions 5.0, 5.1, and 6.0, and also includes an introduction to a few version 6.0 enhancements. This article is not intended to be used as a reference for all the details of WebSphere Application Server administration, but it does describe the key concepts used, and contains a list of references. Although the beginning of the article reviews some fairly basic base server and managed server concepts and operations, much of the remainder of the article will discuss certain complex concepts or operational considerations that will be new even to very experienced enterprise application server administrators.

Part 1 discusses wsadmin deployment to base and managed servers. It also discusses why phased deployments are needed to maintain applications in a WebSphere Application Server Network-Deployment managed cell, and how to maintain high availability in such an environment.

Part 2 of the article will discuss pre- and post-deployment validation and gradual deployment of incompatible versions of applications. It will also discuss the design and implementation of a downloadable Automated Deployment example program that illustrates how to automate the deployment of randomly built collections of enterprise applications or updates, and how to automatically target those applications or updates to the correct servers, including stage-specific application setup.

The Problem

It is quite easy to deploy (install) an application into a WebSphere Application Server setup – typically it just takes a one-line command. For your local running base server, it can be as simple as:

```
wsadmin -c "$AdminApp install X:/
MyApp.ear" -lang jacl
```

However, this simplicity is deceiving and the preceding example is really just the "Hello World" of deployment – a nice demo, but not typical of the real world.

In a real enterprise environment, there are hundreds of interrelated applications spread over dozens of remote application servers, and regular updates that need to be deployed to the right servers, all the while maintaining application availability to users. Even worse, most large enterprises have different sets of operating environments, or stages, each requiring different setups for the same application. For example, the security role mappings and the database used for a specific application in a Microsoft Windows integration stage are likely different from those used for that same application on a Linux production server.

The result might be that at 3 a.m. when some random group of 20 applications has just been rebuilt because of an automated production build of Source Code Management (SCM) code changes, those particular 20 applications each needs to have its updates deployed to its correct individual application server somewhere in the enterprise. And, although it is 3 a.m. in North America, it is prime time elsewhere in the world, so the application updates need to be done in a way that maintains high application availability. This update build and deployment process is regularly repeated, each time involving a randomly different set of updated applications.

Command Line wsadmin and JACL/Jython Scripts

WebSphere Application Server has an extensive administration program, the Administration Console. It also has an equivalent command-line tool, wsadmin, which can be run interactively and which accepts a file of scripted commands. For scripting, wsadmin supports the two script languages JACL and Jython. (See the reference section for links to related material.) In the examples in this article, and in the downloadable example program, JACL is used since it runs on version 5.0 and above (Jython runs on version 5.1 and above).

This article is not intended as a reference for WebSphere administration or the wsadmin tool. For more detailed information, consult the material in the reference section.

Local Base Server Deployment

The simplest deployment scenario is to install an application from a build machine to a local running base server (Network Deployment managed cells are discussed a little later in this article). The local WebSphere Application Server installation has a bin directory containing the wsadmin command and other tools.

The typical sequence for an initial application installation is:

1. If the local server is not already running, then start it:

```
startServer server1
```

2. Install the application:

```
wsadmin -c "$AdminApp install
  C:/MyApp.ear" -lang jacl
```

3. Optionally list all current installed applications (to verify that it really was installed):

```
wsadmin -c "$AdminApp list" -
  lang jacl
```

4. Save this new server configuration:

```
wsadmin -c "$AdminConfig save"
  -lang jacl
```

5. Start the application (this is a one-line command):

```
wsadmin -c "$AdminControl
  invoke
    [$Admi nControl queryNames
    type=ApplicationManager,*]
startApplication MyAppName" -
  lang jacl
```

The typical sequence for an uninstallation using a script file instead of interactive commands is:

```
wsadmin –f uninstall.jacl –
  lang jacl
```

where the file uninstall.jacl contains the following lines:

```
$AdminApp list
set appMgr [$AdminControl
  queryNames type=Application
  Manager,*]
$AdminControl invoke $appMgr
  stopApplication MyAppName
$AdminApp uninstall MyAppName
$AdminApp list
$AdminConfig save
```

Note that the default wsadmin server connection type SOAP requires a running server. It is possible to connect using connection type NONE, but the available operations are restricted. Refer to wsadmin documentation for more details.

Remote Base Server Deployment

A slightly more typical environment is one where one or more target base servers are on remote machines. The most obvious approach (and the one used by some customers) is to use File Transfer Protocol (FTP) to transfer the application or update from their build machine to the remote target machines, and then to use telnet, or some other similar program, to run the wsadmin program on those remote machines. This works, but it is quite messy. It requires FTP and telnet on each remote machine and quite a few error-prone manual operations. It also requires FTP and telnet accounts on each remote machine, which introduces administration and security issues many companies would prefer to avoid.

WebSphere Application Server provides a nice solution to this scenario that many users are not even aware exists. If you install WebSphere on the build machine, then the WebSphere Application Server runtime and its wsadmin command tool are available even if that installation server is never configured or started. You can then use the wsadmin command to connect to your remote target servers and run your deployment commands against those remote servers. If a local build machine file is being installed on the remote server, WebSphere Application Server will internally do the file transfer for you. All the normal WebSphere security (if configured) is automatically used according to the userID and password used with the wsadmin tool. Thus, deployment to remote servers is as easy as to a local server; the only difference is that the wsadmin invocation command specifies a remote server instead of a local one:

```
Wsadmin -host MyRemoteHost -port
  MyRemotePort …commands…
```

Note: The above solution for remote base server deployment only works if both the remote and local installation (where wsadmin is running) are WebSphere Application Server version 6 installations. If either is a version 5.0 or version 5.1 installation, then you will get an error "X:\MyTEMP\appnnnnn.ear does not exist for installation." Base Server (nonmanaged server) Remote File Transfer support is a version-6 enhancement – unfortunately, the version-5x error is not clear about that.

Ellen Matheson McKay is an information developer for IBM Canada Ltd. She writes online help and publications for Rational Application Developer.
ecmckay@ca.ibm.com

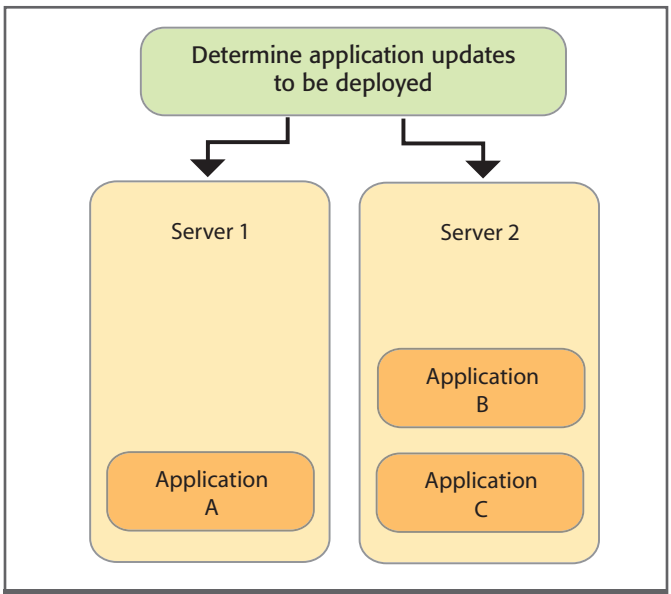

FIG 1: BASE SERVERS

Figure 1 shows a very simple, but representative, base server organization.

This approach provides a nice solution for a build machine and one or two remote base servers. But as the number of independent base servers increases, it rapidly becomes desirable to have centralized administration to manage the collection of servers as a single administrative cell.

Network Deployment Managed Cells, Nodes, and Clusters

As the number of applications and the number of target servers grows, it becomes almost essential to provide single administrative control for these large collections.

A node is a physical collection of one or many application servers on a particular machine, and typically there is one Node Agent per node controlling those servers on that remote machine. Why should there be more than one server per machine? In addition to the obvious redundancy consideration, there are technical and performance reasons (such as Java threading limitations, Java memory garbage collection considerations, etc.) why splitting one larger server into multiple servers can provide significant performance improvements on that same machine. This is particularly common on very large and highly reliable server platforms. For details, see the WebSphere Application Server performance "Best Practices" documents. Note that there can be two or more logically separate nodes (each with its own servers) on the same physical machine, but that is very unusual.

Why should there be more than one target machine? In addition to the obvious ability to supply additional resources (additional CPU cycles, physical memory, etc.) and to provide physical redundancy against hardware failure, there may be organizational or geographical requirements. As well, the various Quality Assurance (QA) environments, such as development, integration, test, pilot-production, and production are almost always run on different machines for logistic, administrative, and security reasons.

To simplify the setup and ongoing administration of redundant servers, most application server products introduce the concept of clusters. A cluster is a logical collection of servers each, of which typically contains the same set of applications. (Any one cluster member could actually contain additional or different programs, but that is highly unusual and confusing.) A vertical cluster contains multiple cluster members on the same node, while a horizontal cluster contains cluster members on different nodes (which is more typical).

A managed "cell" is an administrative collection of many servers. One of the key features of WebSphere Application Server Network Deployment is that groups of servers can be federated together into a single managed cell. A special server called the Deployment Manager (DMgr) manages all the servers and applications in the cell using one Node Agent (NA) per node to control the one or more servers in that node. Node Agents are typically configured as always-running "daemons" within their host machine operating system, so they are always available. They can start or stop servers on that machine, install or uninstall applications on the servers, and can configure other server settings and control other server functions.

Figure 2 shows a very simple, but representative and scalable cell organization.

Using wsadmin with Managed Servers (-conntype SOAP)

In our earlier discussion about using wsadmin to deploy to a local running base server, you probably assumed (correctly) that the wsadmin program connected to that server to perform its various administrative operations. We also could have started wsadmin specifying the default connection parameter type of -conntype SOAP (or RMI). If the local server is not running, we can start wsadmin with -conntype NONE and then wsadmin will directly manipulate the local server configuration files and their contents. In this case, there is less functionality available and some of the commands (like starting an application) are not available. In the case of remote base servers, we supplied a host and port parameter to specify which running server to connect to, so clearly it is not possible to use -conntype NONE for remote serv-

Nothing beats our racks

Absolutely nothing

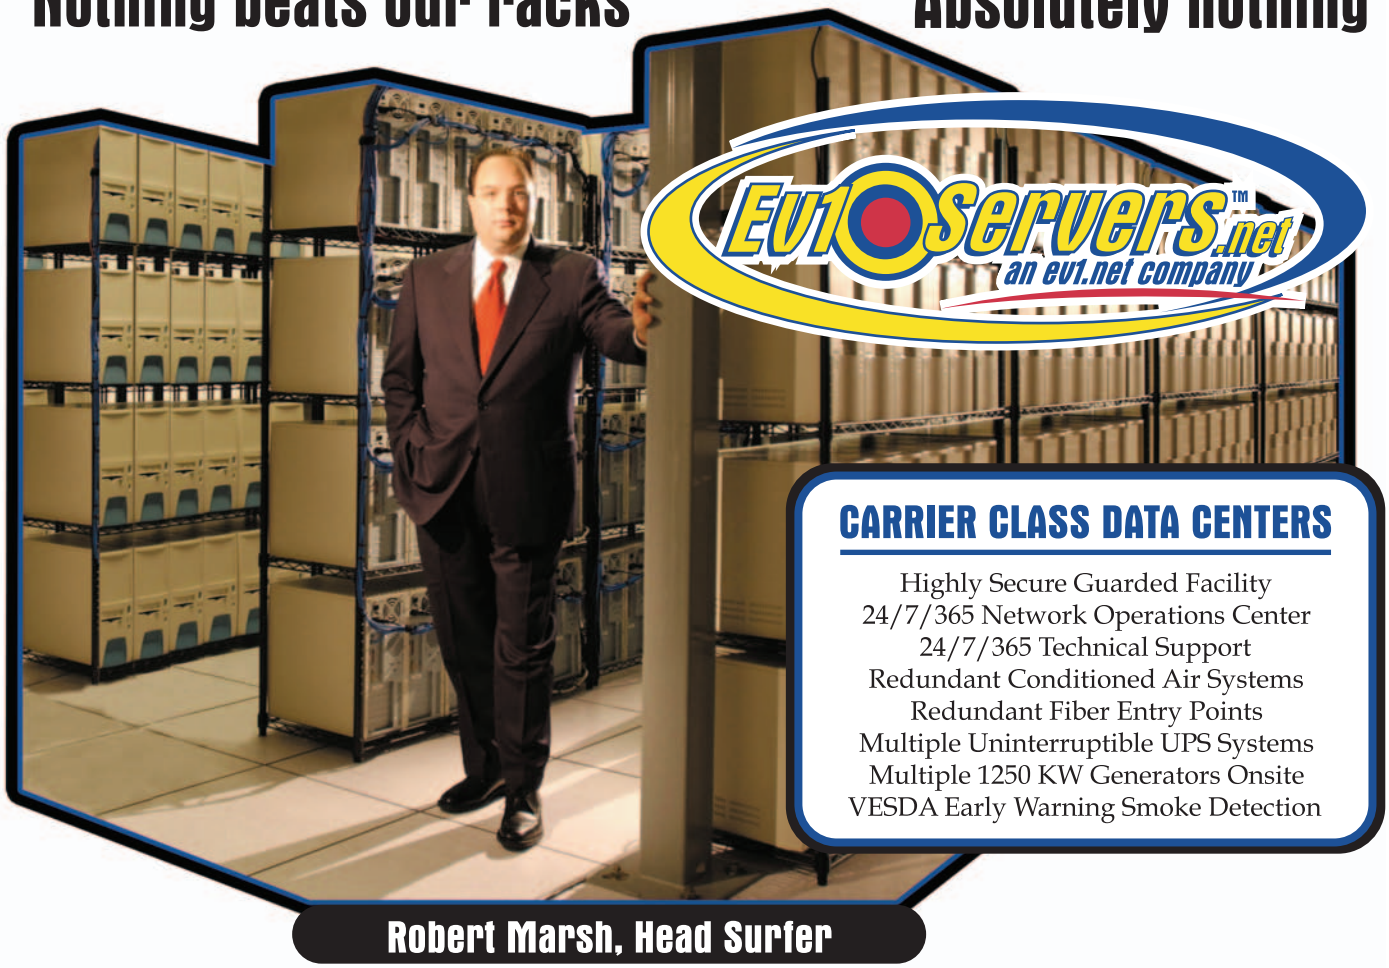

Robert Marsh, Head Surfer

CARRIER CLASS DATA CENTERS

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

START YOUR OWN WEB HOSTING BUSINESS TODAY!

from **$299*** *Instant Activation!*

Dedicated Server

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

Over 20,000 Servers!

1-800-504-SURF | ev1servers.net

PLESK7 RELOADED
Preferred Control Panel

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

ers since the wsadmin program has no direct access to their (remote) configuration files.

In the case of managed servers (servers federated into a Network Deployment cell), determining what is happening is not quite as obvious. If the local server is configured as part of a managed cell, then by default the wsadmin program will connect to its Deployment Manager server and ask that Deployment Manager to perform the specified administrative operations against the specified target servers. Again, the default connection type is -conntype SOAP. For managed servers we can override the default destination and use the host and port parameters to connect to a Node Agent controlling one or more target servers, or to connect directly to a target server. Connecting to the Deployment Manager (the default option) gives the most functionality, to Node Agents a little less functionality, to servers even less functionality, and -conntype NONE to a local server gives the least functionality. It is highly recommended that you always connect to the Deployment Manager, otherwise localized changes might be incompatible with the master cell configuration (such as installing an application with the same name as is already installed elsewhere in the cell) and could cause subsequent synchronization errors. For details, consult the WebSphere Information Center documentation.

How do you have a build machine deploy to remote production servers? Earlier we said that you could have WebSphere Application Server installed on a build machine, without requiring a server to be actually configured or running. Therefore, you can just start wsadmin using the host and port parameters to specify the remote Deployment Manager for the production cell, and everything works as expected. Note that the Deployment Manager must always be at the same version (or a later one) plus service level as all of the servers within the cell.

Note that WebSphere Application Server version 6 provides support for multiple server "profiles" (an expanded implementation of version 5.1 instances), where each server profile is essentially a totally independent server (independent configurations sharing a common set of runtime programs). The wsadmin command will operate against its default server profile (determined by the profile location you are executing from), but you can use the parameter –profileName MyProfile to override that default and specify the actual server profile (and hence its particular Deployment Manager) to be used, or you can use the host and port parameters to directly specify the destination (typically a Deployment Manager).

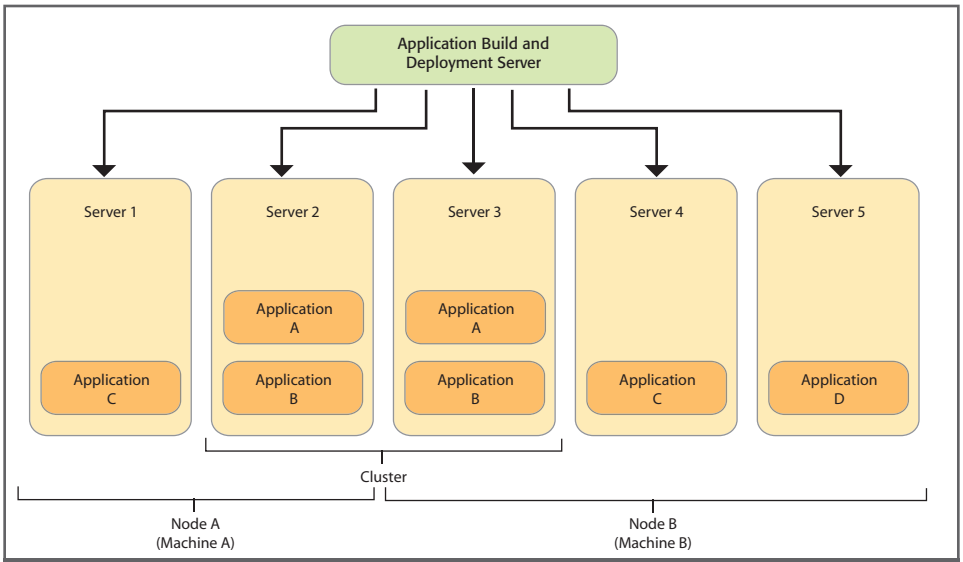

FIG 2: CELL ORGANIZATION (SERVERS, NODES, AND CLUSTERS)

Web Servers, Firewalls, Redundancy, and Workload Management

Note that there are typically one or more Web servers in front of the set of application servers. The Web servers accept incoming user HTTP requests and route each request to an appropriate application server to perform the work. Of course, one or more levels of firewall protection are almost always present at different locations within the system.

Figure 3 shows a simple, but representative HTTP organization.

Good Availability Using Redundancy and Server Failover

Looking after the design, setup, and operation of redundant systems is a significant activity, and there are whole books devoted to this topic. For detailed information, consult the material in the reference section.

If two or more application servers are hosting the same application, then such redundancy can provide workload sharing (better throughput and response time for increasing numbers of requests). This redundancy can also be used to provide reactive server failover and recovery in case of unexpected hardware or software failure. In case of an application server failure, WebSphere Application Server can be configured such that incoming Web server work requests are routed to a different server running that same application. All new requests are handled appropriately, but any

"High application availability can be achieved in spite of unexpected system failures"

in-progress HTTP session will be disrupted and must be reinitiated by users (it will appear to them that there was a brief service failure).

High Availability Using Server Failover and Session Recovery

WebSphere Application Server can reroute new work requests, but it can also be configured to provide reactive session recovery to allow most in-progress requests to continue and complete on the redundant server without any visible interruption to end users. Thus, high application availability can be achieved in spite of unexpected system failures. For most users, even in-progress HTTP sessions appear to continue without any visible interruption.

Continuous Availability Using Preemptive Work Rerouting ("Quiesce")

Even though many HTTP sessions can be recovered, applications with very large session data or other session data that cannot be persisted, can cause an in-progress HTTP session to fail. Thus, if an application or system has a planned outage (due to maintenance or other operator-initiated situations), then preemptive work rerouting can be used to provide nearly continuous availability.

The affected servers are first "quiesced," meaning that all new incoming work requests are routed instead to other redundant servers, and all in-progress HTTP sessions on the quiesced servers are temporarily allowed to run until completed. After a reasonable period of time the quiesced servers are stopped and the planned maintenance or other activity is performed. After the planned activity is completed, the servers are "reactivated." New incoming requests can again be routed to them (in addition to the other, still running, redundant servers).

Quiescing entire cells (as discussed a little later in the Gradual Rollout section) is typically done up front at the Network Dispatcher (IP Sprayer) level using products such as WebSphere Load Balancer. The Network Dispatcher must be configured with session affinity to ensure that all user requests from the same user HTTP session get routed to the same processing organization. Quiescing individual application servers is typically done at the HTTP Web server level by manipulating the routing table in the plugin-cfg.xml file. This file is checked for changes by the HTTP server once per minute, or at another specified frequency. The HTTP server must be similarly configured with session affinity. Note that WebSphere Application Server version 5 could generate the plugin-cfg.xml files, but they needed to be manually transferred to the HTTP servers. Version 6 has a new feature to allow its HTTP servers to be federated into a cell, and to transfer the plugin-cfg.xml files.

AutoSync Application Updates Throughout an Enterprise Cell

WebSphere Application Server has the ability to provide automated distribution of application

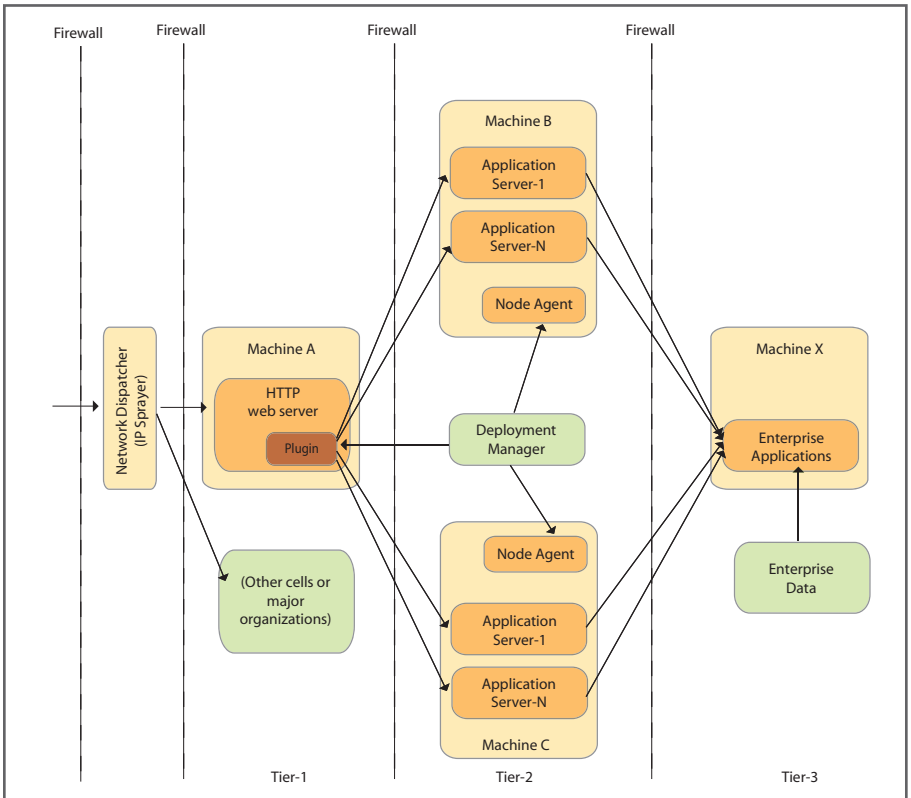

FIG 3: HTTP ORGANIZATION

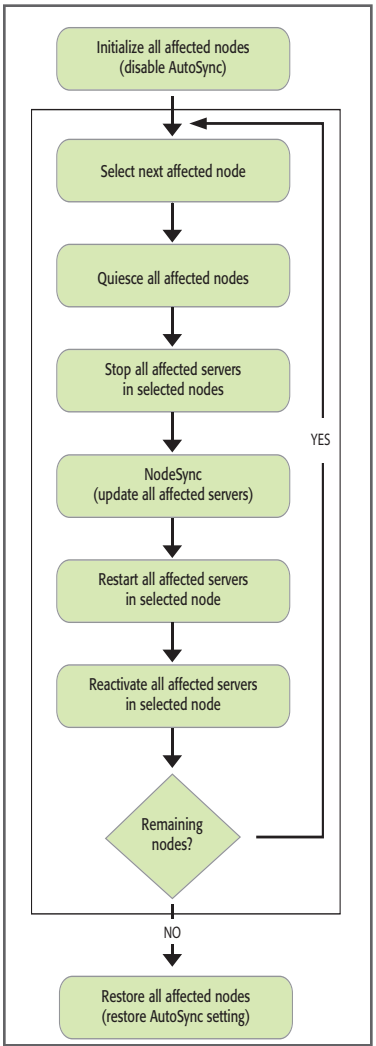

FIG 4: PHASED DEPLOYMENT

updates to servers throughout a managed cell. The application update is first installed into the Deployment Manager application repository. If a Node Agent has its AutoSync feature enabled (which is the default setting), then that Node Agent will periodically (by default, every 60 seconds) perform a NodeSync, which asks the cell Deployment Manager for any application updates. Any such updates are then transferred to the Node Agent, and that Node Agent then updates all its affected servers (since it might have multiple servers that are configured to run the same application). Thus, if each Node Agent has the default AutoSync enabled, any application updated or installed into the Deployment Manager will be automatically distributed to every affected node and to every affected server on those nodes over a relatively short period of time.

Application Updates Can Cause Application Availability Failures

Some administrators fail to understand that default application updates can cause application availability failures even if their entire system is configured for server failover and session recovery. This is because while an application is being updated it is not available to process work requests, and any work routed to it from a Web server will not get a response. Even worse, any other interdependent application will have a service failure if it tries to use the unavailable application instance. This is because the Web servers and the workload management programs don't know that the application is temporarily unavailable. The solution is to first stop the affected application server before doing the application update (not just the application, the complete application server). The stopped application server will be detected, and server failover and session recovery will then take place as expected.

WebSphere Application Server version 6 has the improved ability to perform incremental in-place updates of application components, which will help improve application availability. However, certain types of application component updates will still require stopping and restarting the application, and thus can still result in application availability failures unless the servers are first stopped and server failover and session recovery are active, or unless preemptive work using quiesce and rerouting are performed.

Remember our earlier discussion of the AutoSync feature to automatically distribute application updates to all affected nodes in a cell? There is an availability issue that needs to be understood and handled. Each of the affected applications will be unavailable on its individual servers during the update. Even worse, if several Node Agents happen to request their NodeSync at the same time or at overlapping times, then the affected application on each of those affected nodes will be transferred and updated at essentially the same time, resulting in application unavailability on multiple nodes simultaneously! Where application availability is a serious concern, the solution is to specifically control the phased distribution of updates throughout the enterprise cell.

Better: Phased Distribution of a Single (Compatible) Application Update

As previously mentioned, a specifically controlled phased distribution of an application update (compatible with previous versions) throughout a cell will minimize application availability issues. The required sequence is as follows:

1. Save the current setting and then disable AutoSync on all affected nodes:
   a. Can also optionally save and disable SyncOnStartup.
2. Sequentially, for each affected node, phase-deploy the update:
   a. Select the next affected node to be updated.
   b. Optionally quiesce all of its affected servers (reroute new work requests).
   c. Stop all of its affected servers.
   d. NodeSync that node to retrieve the update and to install it in each affected server.
   Note: Wait to ensure the EAR expansion is complete.
   e. Restart the affected servers.
   Note: Test and wait to ensure the server is running.

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

"The new cluster-update feature is a convenient way to distribute a single update throughout a horizontal cluster"

f. Optionally reactivate the affected servers (to process new work requests).
g. Optionally validate installed application operation.
h. Optionally request manual confirmation to accept (or reject and restore) this update.
i. Repeat for the next affected node.

3. Restore the previous AutoSync settings for all affected nodes:
   a. Including SyncOnStartup if it was optionally disabled.

The above process involves a lot of manual operation and is very error-prone. A script can be created to perform the specific steps, but a different script is required for each different application and each application's different set of associated target servers. The scripts can be quite complex and difficult to create, and require ongoing maintenance as the environment changes.

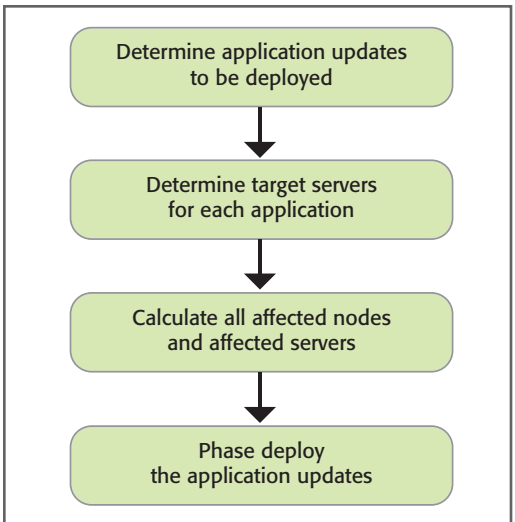

FIG 5: DEPLOYING MULTIPLE UPDATES

There are two special notes in the above steps. First, after performing the NodeSync, the application update (EAR) has been distributed down to the node, but the EAR file must still be expanded into the server installed application directory. Until this EAR expansion is complete, attempting to start the server will produce indeterminate results. An IBM Problem Report has been created about this, and there may, in the future, be a downloadable WebSphere Interim Fix to allow scripts to test for the completion of the EAR expansion. Second, after returning from the wsadmin startServer command, the command has been processed by the Node Agent, but the actual server startup may not yet be complete. Scripts need to test that the server has completed startup and is running.

Note that WebSphere Application Server version 6 has a new cluster-update feature to phase distribute an update to each of a cluster's members one node at a time, including the stopping and restarting of affected servers being updated. Where server failure and session recovery have been configured and provide sufficient functionality, very good high availability will be provided during cluster updates. This feature currently cannot be used to phase distribute an update to unclustered servers and does not do preemptive quiescing or reactivation of work rerouting. If multiple cluster members are on the same node (vertical clustering for scalability), then they will be simultaneously updated (and simultaneously unavailable) as if the Node Agent had performed a regular NodeSync.

Also WebSphere Application Server version 6 cluster updates can only be done for one application at a time. If multiple applications need to be updated within a cluster, then multiple cluster update operations must be individually performed, and each cluster member node and its affected servers will be cycled again during each individual cluster update.

In summary, the WebSphere Application Server version 6 new cluster-update feature is a convenient way to distribute a single update throughout a horizontal cluster with cluster members across multiple nodes, and it does so while maintaining high-availability.

Best: Phased Distribution of Multiple (Compatible) Application Updates

As just mentioned, a specifically controlled phased distribution of an individual application update helps maximize availability. However, if multiple application updates are required, with each application update targeted to a potentially different but overlapping set of enterprise servers, then the process is slightly more complex even though all the same basic principles apply. The sequence required is as follows:

1. Analyze the set of all current updates to be deployed
2. For each application, determine (read) its specific set of target nodes and servers
3. From the total set of affected nodes and servers, calculate the subset of unique affected nodes

and unique affected servers
4. Perform the previously described phased distribution for each affected node

As for the phased deployment of single applications, the above process involves a lot of manual operation and is very error-prone. Additional scripts can be created to assist in these steps, but with ever-increasing numbers of different applications, target servers, and stage environment and applications settings, the number and complexity of specialized scripts becomes an enormous challenge to create and then to maintain.

Part 2: Deployment Validation, Gradual Deployment, Automated Deployment, and References

Part 1 covered the basics of deployment of enterprise application (EAR) updates throughout an enterprise cell. Part 2 will cover pre- and post-deployment validation, and gradual deployment (of incompatible versions of applications). It will also describe the "Automated Deployment" example program in detail, and will include an extensive set of references.

Resources

WebSphere Application Server Administration

- *IBM WebSphere Application Server Network Deployment V5.1: System Administration:* ftp://ftp.software.ibm.com/software/webserver/appserv/library/wasv51nd_admin.pdf *Ch-4: Welcome to Scripting, Ch-5: Deploying and Managing Using Scripting*
- Williamson, Leigh; Chan, Lavena; Cundiff, Roger; (et al.) (2004). ISBN-0131446045. *IBM WebSphere System Administration*
- Barcia, Roland; Hines, Bill; Alcott,Tom; Botzum, Keys (2004). ISBN-0131468626. *IBM WebSphere: Deployment and Advanced Configuration*

WSADMIN Scripting

- *WebSphere Application Server Information Center: Deploying and managing using scripting:* http://publib.boulder.ibm.com/infocenter/wasinfo/index.jsp?topic=/com.ibm.websphere.nd.doc/info/ae/ae/trun_wlm.html
- *WebSphere Application Server Technical Library, Sample scripts:* www.ibm.com/developerworks/websphere/library/samples/SampleScripts.html
- *RedBook: WebSphere Version 5 for z/OS – WSADMIN Primer:* www.ibm.com/support/docview.wss?uid=tss1wp100421

WebSphere Availability and Work-Load Management

- *Extreme Availability with WebSphere and DB2:* www.ibm.com/websphere/developer/zones/hipods/
- *High Performance On Demand Solutions:* www-128.ibm.com/developerworks/websphere/zones/hvws/library.html
- *RedBook: WebSphere V5 Performance, Scalability, and High Availability*: www.redbooks.ibm.com/abstracts/sg246198.html
- *RedPaper: Server Clusters For High Availability in WebSphere Application Server Network Deployment V5:* www.ibm.com/support/docview.wss?uid=swg27002473
- *Maintain continuous availability while updating WebSphere Application Server enterprise applications:* www.ibm.com/developerworks/websphere/techjournal/0412_vansickel/0412_vansickel.html
- *Using WebSphere Application Server V5 for Load balancing and Failover:* www.ibm.com/developerworks/ibm/library/i-wasldbal/index.html

A detailed list of resources pertaining to this article can be found online at http://sys-con.com/websphere.

Alliances and Business Development

An interview with Steve Pattison

STEVE PATTISON

INTERVIEW BY JACK MARTIN

Steve Pattison, vice president of alliances and business development at Interwoven, talked to Jack Martin, editor-in-chief of *WJ*, about the enterprise content management (ECM) provider's relationship with IBM.

Steve Pattison, VP of Alliances and Business Development for Interwoven, oversees all third-party relationships worldwide, including Systems Integrators, technology partners, and OEMs.

WebSphere Journal: All right. What does a typical day look like for you? What do you spend your time doing?
Steve Pattison: I'm responsible for all third-party relationships, including those with systems integrators, and key technology partners, such as IBM.

WJ: You have industrial-strength products and solutions that are well respected in the space and you could have partnered with a multitude of companies. Why did you choose to partner with IBM?
SP: IBM and Interwoven share a common vision for developing applications that are based on J2EE and service-oriented architecture (SOA) standards. IBM is committed to offering extensible and scalable middleware and infrastructure that makes our job of building industry-specific ECM solutions for business a reality for our customers.

WJ: A lot of people partner with IBM and some make a lot of sales and some don't. In your opinion, why is Interwoven doing so well with IBM?
SP: I think a number of factors contribute to our joint success. Interwoven is unique in that we have a native J2EE architecture that allows us to interoperate with WebSphere. As a result, we're able to offer targeted ECM solutions that enable business users to address key content challenges, such as document management and collaboration, productivity, information and process overload, Web change management, brand management, regulatory compliance, and e-mail overload. Additionally, we have a number of OEM VAR customers who build content applications on top of our platform that run on J2EE application servers. Interwoven's business is greatly expanding in the OEM arena and our relationship with IBM has proven to be extremely valuable to that set of OEM VAR customers.

WJ: Why is that?
SP: Customers are looking for superior function and scalability in content applications. For example, OEM customers will use our platform to build content-centric applications like contract management, project management, and client services management solutions. At the same time, these high-function applications are able to leverage the scalability and stability of the IBM WebSphere platform, which is well known to many of our OEM end customers.

WJ: How do you find interacting with WebSphere as a technology?
SP: As a technology we've found the IBM APIs to be very easy to use, especially given the next-generation features of the Interwoven architecture. We've also found that IBM's support organization is consistently there for us in real time, which has been very important as Interwoven continues to push the envelope when it comes to J2EE technology.

WJ: What types of problems do companies have for which Interwoven would be the right choice?
SP: Our customers purchase Interwoven to address specific business challenges that revolve around unstructured content. For example, in areas such as legal departments, where lawyers are working with contracts or other matters, Interwoven offers content applications that map to the unique end-user requirements of legal professionals. So far, the market has innovated greatly when it comes to structured or transactional process areas, such as accounts payable, sales and distribution, and inven-

tory management. However, process areas such as legal, marketing, and engineering project management have historically been underserved. Interwoven is excelling in these types of high-touch content application spaces.

WJ: When a customer is using Interwoven with WebSphere, what is his or her ultimate goal?
SP: Customers want content applications that offer both high function and scale. Working together, Interwoven and IBM help customers achieve this objective by offering industry-specific solutions.

WJ: Where do you see content management going in 2005?
SP: In 2005, the movement is away from back office infrastructure that is sold only to IT, to specific ECM solutions that address the needs of business users across departments, enterprises, and industries.

WJ: Why is that? Why is content management headed in this direction?
SP: The ECM market is headed in this direction because process areas, such as legal, sales, and marketing, have been underserved by technology thus far. In addition, regulatory compliance requirements, such as Sarbanes-Oxley, are mandating that organizations improve and organize key business processes with the support of an underlying content management platform. As our customers strive to find ways to deal with new regulations like Sarbanes-Oxley and Basel II, they need more than just a storage paradigm to resolve these regulatory pressures. Ultimately, what they need are content applications that support the day-to-day work of business users as they correspond with clients, manage key contracts, projects, and associated deliverables. Interwoven excels at offering these deep solutions designed for business users, providing exponentially more than just a file system.

WJ: Why do you excel there? That's a really tough place for companies. Why are you so good at it?
SP: Interwoven built its product and solution suite to address specific end-user requirements. As a result, our applications are very user-friendly, have configurable workspace environments, and require less IT intervention than competing offerings.

WJ: Are you and IBM making joint sales calls?
SP: Yes, we are. Interwoven has a very strong, long-standing relationship with the IBM software group as well as with IBM Global Services.

WJ: A lot of customers read our magazine. If they were interested in finding out more about your product, whom should they call or what should they do?
SP: If they are OEM customers looking to build content applications leveraging our platform, they should contact Interwoven's sales channel, which can offer them the appropriate support and service. If they are customers who are interested in buying our solutions, products or platform, they should contact our worldwide sales organization. We have sales reps in every region across the globe.

"IBM and Interwoven share a common vision for developing applications that are based on J2EE and service-oriented architecture (SOA) standards"

WJ: And if they are IBM customers, could they call their client rep?
SP: Yes.

WJ: When does your product stop scaling, or does it?
SP: Interwoven's platform is designed to scale and support the needs of the world's largest enterprises. Customers can feel confident that Interwoven's solutions running on IBM technology can meet their global requirements in 2005 and beyond.

WJ: When customers consider content management, what should they think about when they think Interwoven?
SP: Customers recognize Interwoven as a leader in Enterprise Content Management. Interwoven offers a complete platform that provides highly functional and scalable Web content management, document management, digital asset management, imaging, records management, and e-mail management capabilities. We leverage this platform to innovate industry-specific solutions that meet the unique requirements of departments and industries. Solutions of this kind bridge the gap between content platform capabilities and what end users require to do their daily jobs more quickly, effectively, and productively for the overall organization.

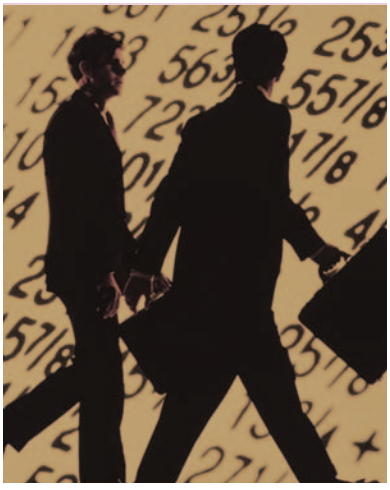

IBM's Middleware solutions save time and money

Telecommunications Companies Face Many Challenges

BY MIKE PERERA

Mike Perera is an IBM market segment manager for Telecommunications
mperera@us.ibm.com

Many of the people using hundreds of cellphone minutes a month, clicking camera phones, downloading huge files on rapid-fire Internet connections, and taking their WiFi laptops everywhere don't remember when telecommunications services in the U.S. were provided almost entirely by AT&T.

But in the more than 20 years since what was called "the phone company" was broken up, the deregulation of the market has propelled the industry into a state of hypercompetition. The coupling of that unleashed competition with galloping technology advances has made it very difficult for telecom companies to be financially successful today.

Furthermore, the three types of telephony services providers – traditional wireline, wireless, and cable – are competing for the same customers. They all face common and unique challenges.

The traditional wireline companies, such as AT&T and the Regional Bell Operating Companies (RBOCs), have invested a lot in their physical networks and have entered other markets including high-speed data (DSL) and wireless, in some cases through partnerships. But customers change carriers frequently and in increasing numbers they are dropping wireline service entirely in favor of wireless. This migration has fueled fierce price competition between wireless and wireline providers, and squeezed margins for both. Many phone companies find their own divisions fighting for the same customers.

Landline companies are also losing revenue as dial-up Internet access decreases and the number of Internet-based phone providers grows, creating more turnover and forcing lower fixed prices for calls. In the U.S., incumbent Internet Service Providers are expected to lose 10.2 million basic access lines by 2008. Additionally, an expected 30% of cable modem users are expected to adopt Voice over Internet Protocol (VoIP) for phone calling. VoIP through WiFi networks in public sites is also taking a toll.

Wireless companies also face intense profit pressure. Industry analysts predict that over the next decade mobile phones will carry 50–80% of all voice minutes, versus today's average of 20%. This has ignited competition from numerous wireless companies large and small, domestic and international. The heated battle to win new customers and keep the existing ones has increased customer acquisition costs while lowering the average lifetime value of customers, ultimately slashing profits. As a result, wireless companies are under continual pressure to provide services that set them apart and encourage customer loyalty.

Some industry observers believe cable companies are best positioned to win the most business in the long run. They have already successfully expanded their television business by offering high-speed broadband Internet access. They're now attacking the traditional telephony carriers by offering VoIP, completing the so called "triple play" of voice, data, and video. Their biggest challenge is how to penetrate the profitable business segment and address their customers' wireless requirements.

Faced with these competitive and technological challenges, telecom companies need to cut costs, especially in their complex infrastructures and in systems that are often inefficient and redundant as a result of mergers and government requirements. They must also improve upon the worst customer satisfaction record of any industry other than the automotive industry.

Furthermore, they must increase revenue per customer by delivering new products that set them apart, with enough lead time to extract profits before the service is commoditized by copycat competitors. They must balance spending on next-generation technologies with the need to maintain existing Public Switched Telephony Network (PSTN), service levels, and cash flow, and still get a sufficient return on investment (ROI).

"Telecom companies need to take the complexity out of their operational (network) and business infrastructure and out of their processes and systems, so they can respond to customer needs and competitive threats at the breakneck speed that is required," said Michael Perera, Software Group Telecom market segment manager for IBM.

Leveraging efficient IT systems is a critical part of telecom companies' efforts to cut costs and get new profitable services to market quickly. They need to integrate their business and IT processes within and beyond their own organizations, allowing them to respond with speed to customers' needs, technology issues, market opportunities, and competitive threats.

IBM is working with telecom companies to supply offerings that provide these capabilities. In fact, 14 of the world's 15 largest telecommunications service providers use IBM middleware to run their applications. IBM is active in industry organizations like the Parlay Group, TeleManagement Forum, and the Open Mobile Alliance in an effort to create best-of-breed solutions and standards.

This experience has been key in the creation of IBM's Industry Middleware Solutions for Telecommunications, which are designed to address the industry's most pressing challenges. The solutions are part of IBM's effort to deliver middleware based on customer preference for buying solutions designed for their industry. Each solution draws on the appropriate capabilities from IBM's WebSphere, Lotus, Tivoli, DB2, and Rational middleware brands, and IBM's industry-specific middleware combined with applications from independent software vendors (ISVs) and industry-expert services.

"The solutions are based on open standards and industry initiatives to promote best practices," said Randy Newell, program director, Software Group Industry Solutions Marketing, Communications & Public Sectors, IBM. "They are designed to make an extremely complex environment more flexible, resilient, and responsive."

The industry standards and best practices include the TeleManagement Forum's Enhanced Telecommunications Operations Model (eTOM) to speed implementation, lower risk, and improve ROI. "eTOM spells out the best-of-breed industry process model and has evolved over the years to reflect the changing needs of telephony service providers," said Randy.

Based on its experience and expertise in the telecommunications industry, IBM has developed the Telecommunications Component Business Model (CBM). This is an industry-specific logical representation or map of the business that reveals its essential building blocks. This model allows the enterprise to be more adaptive to changing customer needs, to prepare for on demand initiatives, and to focus on competitive differentiation. It facilitates the identification of processes that need to be changed to achieve the business goals.

The eTOM Business Process Framework and its associated business process modeling describe the processes and points of interconnection that make up the end-to-end customer operations process flows for fulfillment, assurance, and billing within operations, and for strategy, infrastructure, and product. The Capability Maturity Model (CMM) provides a way to assign a maturity rating to each process. CBM and CMM provide a logical template for process and capability analysis and process identification for the eTOM Business Process Framework.

IBM also has a suite of products that helps translate the business requirement into technical reality. WBI modeler provides a way to model the process in its current state and use input from the business to convert it for the future. In many cases, the future model is eTOM's industry best practice for which IBM has developed WBI Modeler process models. Once the future state of the model is agreed upon, it is translated into Business Process Execution Language (BPEL) which can be imported into WBI as draft development code. The Rational Unified Process provides a platform to iteratively mature initial development code. The solution has proven very compelling with telecommunications customers around the world as the lines between the responsibilities of Line of Business, IT, and the telephony network become more closely integrated.

"Many phone companies find their own divisions fighting for the same customers"

The IBM solutions for telecommunications are also designed within IBM's industry-leading Service Provider Delivery Environment (SPDE) architectural framework. SPDE enables telecom service providers to bring new revenue-generating services to market faster, easier, and at a lower cost by providing a flexible framework that addresses a service provider's entire environment.

SPDE supports multiple components, including third-party services, applications, and content. It is supported by a community of over 200

ISVs that work together to create comprehensive solutions based on SPDE and include IBM products and services. The result is faster deployment of a wider variety of products and services.

These new products and services include converged services, such as multimedia messaging for consumers, digital video conferencing for families and businesses, and location-based services such as helping travelers locate a specific type of restaurant nearby, or helping an appliance repair company find its nearest truck to a customer call. SPDE also allows service providers to deliver personalized services and offerings based on customer preferences and tendencies.

These kinds of services require systems that are integrated across information and business processes. The IBM Middleware Solution for Telecommunications for Operational and Business Support Systems (OSS/BSS) Optimization can help bring this about. It enables service providers to integrate processes across disparate systems in a cost-effective, flexible manner while continuously optimizing thousands of processes and the corresponding IT infrastructure.

This integration of business and IT helps lower operational costs by improving company-wide efficiency and flexibility and maximizing the use of systems. It also helps increase revenue by getting new customers on board faster, and speeding new products to market.

Specifically, the integration of processes such as order handling, billing management, fault resolution, and customer problem management helps reduce the costly churn of customers that all telecom companies face. Furthermore, the creation of a comprehensive enterprise platform creates a single view of customers, which improves customer satisfaction and allows companies to rapidly respond to customer and market demands.

OSS/BSS is used by 16 of the world's largest services providers, including Nextel, Mobifon, Bharti, Telekom Austria, and Verizon.

Service providers also need to reduce the complexity and cost of their customer service processes. Not only are those overall costs high due to high customer churn or turnover, but the average customer service cost is rising. To compete, providers need to give customers access to what they want, when they want it, and how they want it – in short, consistent, strong customer service across all contact channels.

The IBM Middleware Solution for Telecommunications Contact Center Optimizer provides the tools needed to manage customer care efficiently to lower costs and improve customer satisfaction. It does so by leveraging information in providers' existing systems to present a single view of the customer and provide insight into their value, tastes, and preferences. As a result, customers can receive personalized support via a range of devices (phone, Web, etc.), can manage their accounts themselves, and can be offered targeted promotions to meet their needs. The solution can reduce customer support costs by as much as 40%, while also increasing customer satisfaction and average revenue per user (ARPU).

In addition to cutting costs, telecom companies need to increase revenue by efficiently delivering new offerings that set them apart from competitors. This is critical as providers grapple with their dwindling revenue sources and profit margins. These new offerings include mobile services; products and services from partners including digital content for business applications (CRM, order management) and consumer applications (ring tones, games, multimedia); and next-generation network services such as enhanced voice and multimedia conferencing.

Enterprise customers and consumers are asking for these services and most want them from a single vendor. Product introduction cycles have been long and costly, however, and some new offerings are dependent on the upgrading of existing network technologies.

The IBM Middleware Solution for Mobile Services Delivery helps simplify the process of securely delivering new services. It does so by leveraging existing data and network services, supporting development by internal and external teams across geographies, integrating internal processes with those of third-party digital content providers, providing secure single sign-on for customers and partners, and supporting mobile services for workforces and customers.

The results can include faster introduction of time-sensitive content, more cost-effective development and deployment of mobile technologies, and the introduction of new revenue-generating location-based services, all of which help service providers be more competitive and ultimately deliver higher customer and shareholder value.

"Wireless companies are under continual pressure to provide services that set them apart and encourage customer loyalty"

Ignorance isn't bliss.

Knowing about performance problems before they happen is.

It might seem like bliss to ignore application performance problems and deal with them only if and when they surface. But waiting until a problem arises in production only results in crises, stress, and wasted time. Experience real bliss by eliminating performance problems proactively with DiagnoSys.

From development through production, DiagnoSys targets potential problems by providing a wide breadth of information, collecting data not only from the application server but also the associated resources, such as the Web server, database, and network.

This breadth of data gives a more accurate picture of a developing problem. Other products only look at the app server and imply performance of the resources based on the performance of the connections, a method that often misses the real source of the problem.

DiagnoSys also provides in-depth, statistically correlated information so you see a more complete view. DiagnoSys gathers information down to the method on the app server as well as in-depth information for all resources down to the OS.

Even though you get depth of information, you won't be overwhelmed. DiagnoSys synthesizes the data and quickly identifies the root cause so you can correct the problem before it impacts you or your company.

diagnoSys

Improve Web app performance at your pace. Listen to expert Web seminars – it's a fast and easy way to keep up on the latest advice, tips, and best practices.

Visit: www.hwcs.com/wsj19.asp

SILVER ANNIVERSARY 1979 - 2004

H&W

www.hwcs.com | 1.800.338.6692

Bridging the Enterprise Computing Gap For 25 Years

Achieving accessibility in WebSphere applications

Make All Available to All

BY: MAX KING

Max King is based in Prolifics' London office and has extensive experience delivering WebSphere solutions for clients worldwide, including JP Morgan Chase, NYSE, Wal-Mart, Lufthansa, BNP Paribas, MetLife, Honda, and Hertz. As a senior member of the Prolifics' WebSphere Consulting Division, he specializes in J2EE architecture and best practices, application development and deployment, production crit-sit support, production readiness assessments, application migration, and messaging and integration. He also holds certifications across all WebSphere versions, UML, XML, and Project Management. Max was part of the IBM/SIAC team which was awarded the IBM STAR Award for excellence in teamwork and delivery.
mking@prolifics.com

As your company makes more of its applications available on the Web, you will need to determine the level of accessibility for each of those applications. The term accessibility describes how IT hardware, software, and services address and/or neglect the needs of a user community, including users with disabilities. To achieve this aim, accessible applications often include or interface with assistive technology such as screen reader software, voice recognition, screen magnifiers, and special keyboards.

For WebSphere Web- and applet-based applications, accessibility compliance is primarily a consideration of user interface design:

- Information should not be lost when using an alternative presentation, such as with assistive technology (e.g., processing images with text-to-speech software).
- Site navigation and input controls should also be available in non-default usage scenarios (such as keyboard-only access).
- The overall user experience should not be a diminutive one for users with disabilities – resolve problems of information loss, inaccessible controls, and layouts that may be frustrating or long-winded from different views.

Users today expect higher levels of convenience and ease of use from software applications. However, wider adoption of the technology by new demographic groups, such as older users and special-needs communities, leads to new requirements of the software for accessibility. These requirements go beyond enhancements of convenience and constitute functional changes. Software compliance may address the needs of those with limitations of dexterity or mobility, vision impairments, loss of hearing, eye/hand coordination problems, and cognitive disability. Consider that images flashing at a certain frequency can incite epileptic seizures in some users; this is an example of why accessibility factors need to be taken seriously.

Some of the measures that address special-needs users include:

- Text-to-speech synthesizer software for visually impaired users
- Greyscale mode and high-contrast settings for color-blind users
- Keyboard-only alternative for users with limited hand usage

These few examples of functionality are detailed in government standards for accessibility. The most commonly referenced standard is Section 508 of the U.S. Rehabilitation Act, which places a requirement on all U.S. federal agencies to use products and services that are accessible to people with disabilities. Similar directives exist in the EU, Japan, Australia, and across the globe.

Assistive Technology

Assistive technology includes software and hardware solutions such as screen reader software, voice recognition, screen magnifiers, special keyboards, and wireless communications. IBM provides screen reader software called Homepage Reader which is available as a beta download. Another example of assistive technology is the MS-Windows NT/XP operating system Accessibility Options (see Figure 1). These are reached via:Start Menu -> Settings -> Control Panel -> Accessibility Options.

These options include:

- Settings for cursor flashing
- The enabling of special keyboard settings such as StickyKeys, Filter Keys, ToggleKeys
- The settings for high-contrast mode, including different larger and bolder fonts, combined with customizable background/foreground color schemes

Figures 2 and 3 of the same Windows XP desktop illustrate the effect of high-contrast settings.

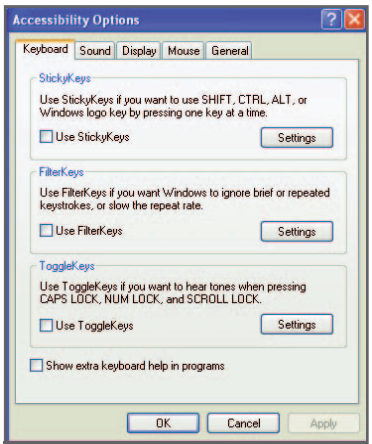

FIG 1: ACCESSIBILITY OPTIONS

As shown in Figure 3, the application's background has become black and many of the fonts are now bolder and larger. The font settings are adopted from the new system settings. Note that the figures include a Web page in Internet Explorer 6.0 and windows from MS Outlook and MS Word applications. All of these applications converted properly to the high-contrast settings, but this is not guaranteed to be the case.

Readers are encouraged to try out these settings and explore how different software packages and applications are affected. The results will vary according to software and version, and depend on whether or not it ignores or blocks Windows systems settings or overloads functions such as special keys. In general, the software needs to support these APIs explicitly, so it is not safe to assume it will be accessible by default.

Web Design Considerations: HTML

The main focus of accessibility considerations for Web applications is on how the content is presented using HTML, stylesheets, JavaScript, and media content. This is irrespective of whether the rendered pages are static or are created using dynamic server-side technology such as Servlets, JSPs, or Windows ASPs. Since it is the final product rather than the creation process that is key, the principles will apply to different types of architectures. This article focuses on WebSphere applications.

Clean Up the HTML

Many browsers are flexible in accepting HTML that does not meet HTML and XML standards. You cannot assume your rendered HTML meets specification even if it looks acceptable in the predominant browsers such as Internet Explorer and Netscape. The HTML needs to be cleaned up; assistive HTML readers do more than handle basic presentation and therefore require a coherent base to interpret. They are not as flexible in handling unstructured HTML as are normal browsers.

Here are some areas that will need to be addressed:

- ***Add missing HTML end tags***. Single tags such as <br> also require an end tag, as per XML guidelines.

BEFORE	AFTER
<td>Hello</td>	<td>Hello</td>

- ***Avoid uppercase HTML tags.*** XML is case sensitive so use lowercase tags only.

BEFORE	AFTER
<TD><B>Hello</B></TD>	<td>Hello</td>
<TABLE>	<table>

- ***Use quotation marks in tag parameters.***

BEFORE	AFTER
<td bgcolor=ffffff> Hello</td>	<td bgcolor='ffffff' >Hello</td>
<table border=0 width=100%>	<table border='0' width='100%'>

Such changes will appear trivial at first, but if these formatting issues permeate most of the site's pages, a seemingly simple task becomes a large and tedious project.

There are tools to assist with the HTML cleanup work that automate fixing problems such as those highlighted above. IBM WebSphere Studio Application Developer 5.1 provides a facility to clean up HTML and JSPs. The "Cleanup" function is available via Web Perspective -> select HTML/JSP > right-click Cleanup Document (see Figure 4).

Clean Up the Forms

Before using automated tools to clean up the HTML, first check the format of forms. This requires a manual process, otherwise, if the HTML is reorganized by a tool (e.g., WebSphere Studio -> Cleanup Documents -> Insert Missing Tags) and form tags are moved around, you will likely lose information that is intended for submission to the server, which will break the application's functionality.

For example, an existing form tag might have been inserted between a table tag and a table row <tr> tag. Although this is structurally wrong and should be avoided, it still works in common browsers, so you might already be using it.

```
<table>
<form name="form1">
  <tr>
     <td><input type="text"
      name="text1" /><input
      type="submit" name="sub1" />
  </tr>
  </form>
</table>
```

Experience has shown that cleanup tools often take liberties when reorganizing form tags, effectively removing tag information that you expect to be part of the form:

```
<table>
  <tr>
     <td><form name="form1"></
      form><input type="text"
      name="text1" /><input
      type="submit" name="sub1" />
  </tr>
</table>
```

This will render a button that does not react when clicked. It is recommended that you move the <form> tags outside the <table> tag altogether. This example also illustrates why you should conduct a functional regression test after making accessibility changes, even though they are largely considered to be "cosmetic" rather than functional.

Another area within forms to consider is the use of labels. Labels are generally

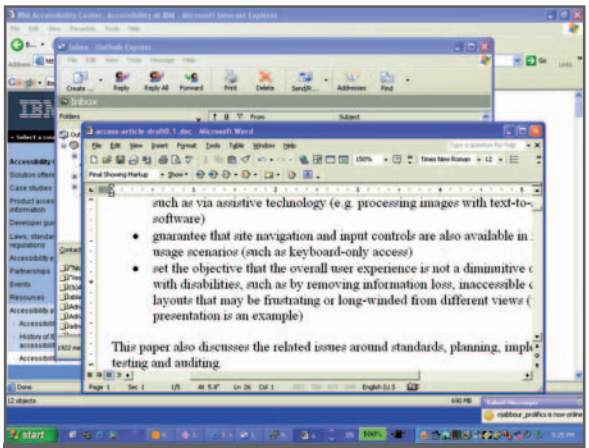

FIG 2: THE EFFECTS OF HIGH-CONTRAST SETTINGS

located close to the entry field (on top of, or to the immediate left of), but that does not guarantee assistive technology will associate it with a particular form element. Use the <label for="id"> element to specify the association between the label and the entry field.

Images Need Alternative Text

Images need text alternatives in the form of ALT attribute or LONGDESC attribute (for longer text) to describe their content in order for reader software to convey their message to users. This is particularly important if the image is a diagram with information content such as figures, pie-charts, etc. That information must be conveyed in a different format as well.

FIG 3: ANOTHER VIEW OF HIGH-CONTRAST SETTINGS

Image Maps Need Alternative Text

Consider a dynamic image map such as a U.S. map of all 50 states in which you can access information for any state by clicking on, or hovering over the state (representing an image hot spot) with the mouse (see Figure 5).

An accessible version of this would need to convey the same level of information for the top-level image, such as the color schemes and what they indicate. It must also address how to navigate to each state via the keyboard.

Accessible Alternatives for JavaScript

The preceeding example included hover-over popups for image maps. As hover-over popups, info boxes, and flyover menus become more common on Web sites, care must be taken to ensure that there is an accessible equivalent for the navigation and information delivery functionality. Flyovers are usually implemented using JavaScript which is not accessible by default. Turn off JavaScript and test the page again to get a better feel for the dependencies, and then limit these as much as possible. Also test the page's functionality in keyboard-only mode.

Easy Navigation and Shortcuts

Many Web sites include a header and left-navigation section, allowing users to use links to jump directly to the page contents section.

The jump links can be achieved by creating a hyperlink and associating it with a one-pixel sized image that has alternative text.

```
<a href="#main"><img src="pixel.
gif" height="1" width="1" bor-
der="0" alt="Jump directly to main
section" /></a>
…
<!–jump to here –->
<a name="main"></a>
```

In case a Web site is using frames to differentiate the screen sections, each frame should have a title in order to improve the ease of navigation among them.

Alternatives for Multimedia Content

Multimedia content generally requires third-party utilities and is rarely accessible; therefore, you must provide accessible enhancements such as captions, visual queues for audio alerts, and volume controls. If enhancements such as these are not possible, you must provide an accessible alternative, for example, a text-only transcript.

Explanatory Text for Web Pages

Consider a Web page that contains a table with sales figures per sales quarter.

In order to understand the numbers presented in any given data field (such as "$160"), the user visually checks the row header on the left, such as "Apples," and then the column header above, such as "2003/Q3." With an HTML reader users cannot perform the same correlations for that data since it is presented in a predominantly linear manner, word for word, as text.

Figure 6 shows how users with unimpaired vision can assume a birds-eye view and have the capability to refocus on special details, shift visual focus to jump around and correlate information, or drill down into other sections. This is very difficult with a reader, therefore special provisions are needed.

A basic technique is to provide explanatory text using a zero-sized image with ALT text. The ALT text will only be presented to the reader software, not to someone viewing the page using a normal browser.

```
<img src="pixel.gif" height="1"
width="1" border="0" alt="Sales of
Apples for Q3/2003" />
```

The example also includes three radio buttons: Details/Edit/Delete. If the radio buttons are not annotated for reader software, the values would be read without the description, such as:

```
Checked
Unchecked
Unchecked
```

More hidden text can be included to yield a less confusing result:

```
[More Details for Apples] Checked
[Edit Apples Numbers] Unchecked
[Delete Apples Data] Unchecked
```

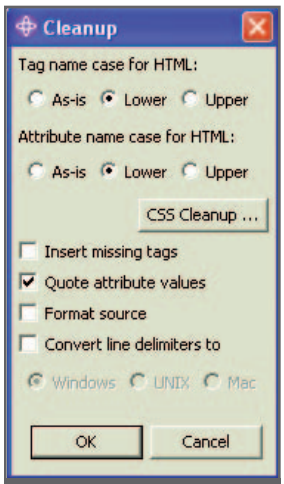

FIG 4: CLEANUP MENU

eclipse CON™ 2005

2nd Annual Gathering of the Eclipse Community

Plug-in to Eclipse

This is your opportunity to get in-depth technical information from the Eclipse experts, learn the latest tips and techniques for using the tools, network with fellow enthusiasts and experience the breadth and depth of the Eclipse community. Attending EclipseCon™ will expand your knowledge and make Eclipse work better for you.


Tim O'Reilly


Urs Hoelzle


John Wiegand


Erich Gamma

Keynotes: Tim O'Reilly,
CEO of O'Reilly Media

Urs Hoelzle,
VP of Engineering, Google

John Wiegand and Erich Gamma,
Eclipse Platform PMC Leaders

Over 60 Sessions and Tutorials for beginner through advanced participants. Some of the conference sessions and tutorials include: Eclipse Platform · Rich Client Platform (RCP) · Test and Performance Tools · Debugging · J2EE Development · Web Development and Tools · Building Plug-ins · Modeling Framework (EMF) · C/C++ Development · Business Intelligence and Reporting · Embedded Development

February 28-March 3, 2005
Burlingame, CA

Register at www.eclipsecon.org

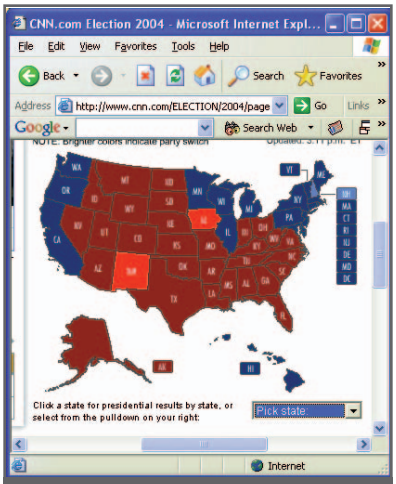

FIG 5: IMAGE MAP FOR CNN ELECTION RESULTS

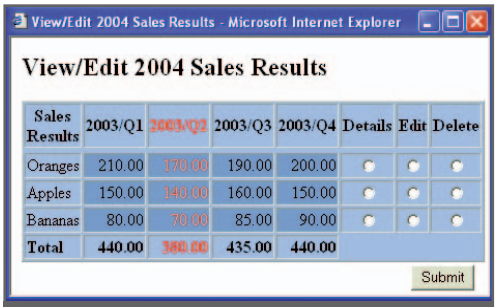

FIG 6: SALES FIGURES PER SALES QUARTER

A more elaborate strategy than hidden text would include using Model/View/Controller components and allowing the Model to be rendered differently in the View according to the client's settings. Portal software such as WebSphere Portal Server can be used to implement this strategy.

Data Tables vs Layout Tables

HTML tables can be used both to store data and organize page layout. It is important to differentiate these usage types. Layout tables should be flagged so they are not interpreted as data tables. One convention is:

```
<table summary="layout">
```

Caution with Color Characteristics

Colors are often used in Web pages to highlight or differentiate text segments, such as red text to point out critical information. However, these colors will not be picked up by reader software, color-blind users, or high-contrast settings.

Use a hidden ALT text image to highlight these characteristics. In the sales table example (see Figure 6), the Q3/2002 column has critically low sales figures highlighted in red, but that information needs to be conveyed to reader software explicitly.

Avoid Popups

Popup windows are not recommended since many clients have activated popup blockers that disable them. Additionally they cause unnecessary disruptions to HTML reader software because they force a context-shift to the new window. Action is then required to close the popup or shift back to the previous window.

Allow Stylesheets to Be Deactivated

What will your Web site look like with CCS stylesheets disabled entirely? Does this lead to information loss?

Rephrase Descriptions for Mouse Usage or Visual Cues

Wording such as "view this" and "click here" will not address all users, so rephrasing these is appropriate.

Web Design Considerations: Java Applets

Web applications are generally HTML-based and can be viewed directly in a Web browser. One other type of application is the Java applet that is loaded over the Web via an HTML page, but provides its own interface. Java Swing classes include an Accessibility package that should be used to set image labels and other accessibility options. Swing's system look and feel should also be used. If the applet (or Java standalone GUI) is still using the AWT classes, strongly consider migrating to Swing.

Other considerations for designing desktop applications as well as Java applets include:

- Alternatives to audio queues such as beeps
- Clear focus on the current part of the screen
- Easy navigation between components via both mouse and keyboard
- Access to the operating system
- Accessibility APIs such as customizable colors and high-contrast settings
- Provide queues such as labels for reader software
- Allow timed areas to extend the time permitted to perform action
- Avoid blinking text or graphics

These considerations also apply to HTML-based Web applications, but are less common.

Conclusion

According to a survey in 1997, over 17% of the adult population in the U.S. had a disability. With corporate applications rapidly getting deployed on the Web, it is critical and at times mandatory to address accessibility for users with limitations of dexterity or mobility, vision impairment, loss of hearing, eye/hand coordination problems, and cognitive disability.

Accessibility cannot be considered an afterthought in Web page design since it touches on so many different areas of applications. For existing Web applications, it is possible to enhance the various accessibility options, but with the number of bulk changes that may be necessary, the effort should incorporate several iterations and must include testing. Testing application accessibility is very different from testing applications; testers need to detect problems in areas that are not visible to most users and that will not be addressed by Web development. Development and testing teams should agree on the scope, since there are potentially many grey areas to cover. The existing standards can help to define this scope, especially if compliance is required.

Useful Links

- *IBM Accessibility Center*: www306.ibm.com/able/
- *Federal Rehabilitation Act (section 508)*: www.section508.gov, www.access-board.gov
- *W3C Web Accessibility Initiative*: www.w3.org/WAI/
- *UK Disability Discrimination Act Code of Practice*: www.drc.org.uk/open4all/law/Code%20of%20Practice.pdf
- *IBM Corporate Instruction 162*: www-306.ibm.com/able/access_ibm/execbrief.html

LinuxWorld® CONFERENCE & EXPO

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

HYNES CONVENTION CENTER • BOSTON, MA

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

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

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

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

LINUXWORLDEXPO.COM

OPEN Where MINDS Meet

Register Online With Priority Code: D2601

PLATINUM SPONSORS

BakBone Redefining Data Protection. · ca Computer Associates · hp invent · IBM · Novell · ORACLE · Sun microsystems · IDG WORLD EXPO

A Mainframe Grid Computing Infrastructure

Based on the IBM Grid Toolbox

BY DIKRAN S. MELIKSETIAN, MOON J. KIM & COLM MALONE

Moon J. Kim is an IBM senior technical staff member and program director responsible for the development of IGA strategic infrastructures and businesses. He has developed products and solutions such as the grid technology, the advanced Web system, and broadband systems and devices. He was a key architect of memory and I/O systems and holds many patents in this area. He is an IBM Master Inventor and has published numberous books and papers.
moonkim@us.ibm.com

In an on-demand business, the computing environment is integrated to allow systems to be seamlessly linked across the enterprise and across its entire range of customers, partners, and suppliers. It uses open standards so that different systems can work together and link with devices and applications across organizational and geographic boundaries. It is virtualized to hide the physical resources from the application level and make the best use of technology resources thereby minimizing complexity for users.

A good example of a virtualized model is a grid computing model. In this article, we describe the design of a grid based on IBM mainframes. The system under consideration is an experimental grid built using the IBM Grid Toolbox version 3 (IGT), which is a commercial derivative of the Globus Toolkit version 3 (GT3). It provides interoperability among zSeries and S/390 processor families and creates an environment that enables the sharing of underutilized computational resources. Typical mainframe systems are designed and provisioned with an average utilization target that always reserves extra capacity for future growth and handles unexpected business-critical job load demands. The unused capacity on these systems is designated as white space and can be leveraged for on demand use.

A number of clustering solutions for zSeries and S/390 systems, such as Sysplex and OS/390 Workload Manager, have been developed to share these underutilized resources. However, these clustering solutions work only among coupled systems that belong to a single family of processors. The zSeries and the earlier S/390 systems belong to multiple generations of processors, such as Generation 5 (G5), Generation 6 (G6), and zServer, and a typical data center includes a combination of those systems. A method to offer interoperability among systems of different generations is needed. The design we offer provides this capability.

Mainframe systems provide a level of isolation and security that is unparalleled on other platforms. For example, the zOS supports logical partitions and guarantees there is no leakage of information or execution resources between them. This allows for the opportunity to provide better isolation between tasks – something not available in grids based on other platforms. The isolation between grid tasks is based on whatever isolation mechanism an operating system on those platforms can provide. Usually, grid tasks are run as separate processes within the operating system, thus sharing resources controlled by the OS. This situation might result in intentional or accidental exposure, or corruption of the data of one task by another task. The design we propose exploits the capability of running multiple concurrent virtual machines on zSeries and S/390 systems to further isolate task execution in separate virtual machines.

Overview of the System

Figure 1 shows an infrastructure built on the Open Grid Services Architecture (OGSA). It offers services for automated setup and management of users via the Web.

As you can see, Figure 1 is a topological view of various data centers geographically dispersed around the world. Grid users interact with the system through a grid management center. The grid management service includes a job management service, as well as general administration and user management services. Each data center includes one or more zSeries or S/390 nodes. The nodes at a given data center may be homogeneous or heterogeneous. In Figure 1, each data center has a combination of different generation machines.

Figure 2 shows the mainframe nodes are partitioned via logical partitioning. Each logical partition, or LPAR, functions as a separate system with a host operating system and one or more applications. Further, each LPAR has one or more logical processors each of which represents all or part of a physical processor allocated to the partition. The logical processors of a partition may be either dedicated to the partition so the underlying processor resource is reserved for that partition, or shared with another partition so the underlying processor resource is potentially available to another partition.

In general, one or more of the LPARs are dedicated

to applications regularly hosted at the data center. However, a special partition can be created for use by the grid infrastructure using the white space. This partition, the Grid LPAR, shares its virtual processors with other partitions at a lower priority than the other LPARs. This ensures that the grid use of the node does not impede the regular operation of the mainframe; it only uses the excess capacity of the mainframe.

Each Grid LPAR runs multiple Linux Virtual Machines (VM). One of those Linux VMs, the Manager Linux VM, acts as a manager and interfaces with the grid management center. The grid management system allocates jobs to the Grid LPARs based on the resources available in that LPAR. The grid management system also migrates jobs from one machine to another when the resources available in a given node are reduced beyond the requirements of the job(s) running on the system.

A Grid Service hosting environment is configured to run in this Manager Linux VM. In our experimental system, we used the OGSI (Open Grid Services Infrastructure) runtime environment provided by the IBM Grid Toolbox. In addition to the standard grid services, a modified version of the Globus Master Managed Job Factory Service is to be deployed. Besides the Manager Linux VM, additional Linux VM instances or Job VMs are configured. The IBM Grid Toolbox with a modified Local Managed Job Factory Service is deployed on these VMs.

The Globus Toolkit and the IBM Grid Toolbox

The Globus Toolkit version 3 (GT3) is an open source, open architecture project that provides a platform for developing grid services and grid applications, as well as a grid service runtime environment based on the GGF (Global Grid Forum) OGSI standard. It also provides a set of tools for a grid administrator to manage grid systems.

IBM Grid Toolbox for Multiplatforms v3.0 provides a fully integrated alternative to the open source distribution of Globus Toolkit 3.0 with additional IBM value-adds. It provides a GGF OGSI grid service runtime environment based on the embedded WebSphere 5.0.2 server. The embedded WebSphere server provides a robust and scalable environment to run grid services. IGT also provides additional features both in usability and services. It uses a simplified wizard-based installation mechanism thereby decreasing the cost of building and deploying grid infrastructures. It also provides an information center and a rich set of tutorials and samples to assist in the design of grid infrastructure and services.

IGT also includes additional grid services, such as an enhanced Registry service, a Common Resource Management (CRM) service, and a Policy Management service. A Web-based management UI (user interface) that facilitates the task of a grid administrator is also included.

The Policy Services in the IGTs enables administrators to define a set of business goals and to enforce a set of rules that allow their grid to meet these goals. In the IBM Grid Toolbox, a policy identifies the desired outcome for the interactions between different elements in the grid environment.

Components

Most of the components within the Manager Linux VM are based on the corresponding IGT/GT3 components and services as shown in Figure 3.

The Manager VM Hosting Environment is the equivalent of the Master Hosting Environment. The Virtual Host Environment Redirector is the same as the Globus Virtual Host Environment Redirector. It accepts all incoming SOAP messages and redirects them to the appropriate Job VM. The Job VM Factory Service is a modification of the Globus Master Managed Job Factory Service. The PortType of this service is an extension of the OGSI Factory PortType. The JobVMFactory PortType has an additional shutdownService operation. The Job VM Factory Service is responsible for exposing the virtual GRAM service to the outside world. It configures the Redirector to direct createService and shutdownService calls sent to it through the Job VM controller. The Job VM Launcher is the equivalent of the Globus Hosting EnvironmentStarter service. Instead of starting a new process, it starts a new Job VM and communicates with that VM.

The grid-mapfile is used to obtain the username corresponding to a particular subject DN (distinguished name). The Job VM Controller ensures that one Job VM is run for each subject DN on a node. When a request to resolve a URL is received, the Job VM

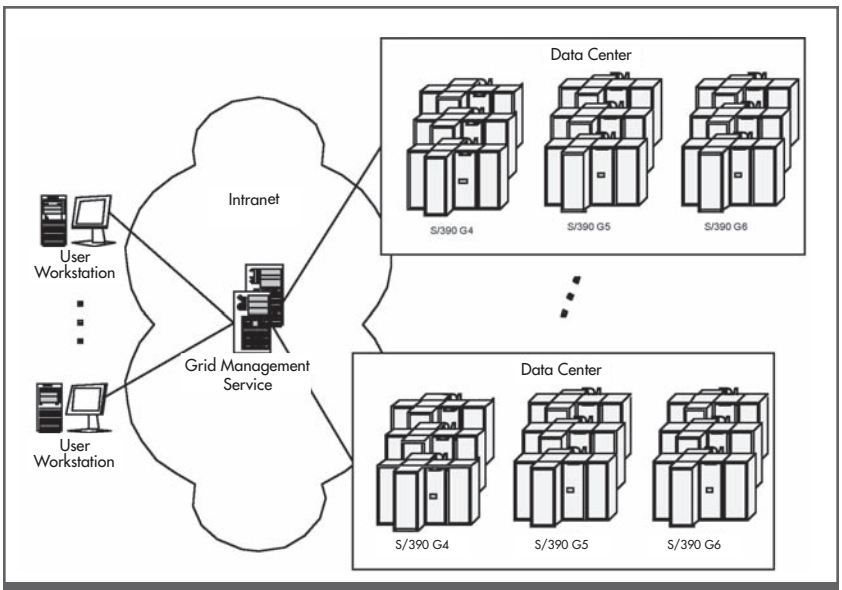

FIG 1: MAINFRAME GRID ARCHITECTURE

Dikran S. Meliksetian is a senior technical staff member with the High Performance On Demand Solutions team in the IBM Software group. He is engaged in the design and development of grids and infrastructure management solutions. dikran_meliksetian@us.ibm.com

Colm Malone is a researcher in the IS department at the Thomas J. Watson Research Center. He has a BS in computer engineering from Trinity College in Dublin. He received a Chairman's Award for developing the Linux Client for e-business (an image used throughout IBM) and transitioning this effort to IBM's Global Services division. Recently his work has involved the area of Grid Computing, specifically leveraging the computing potential of the many Linux platforms within IBM. colm@us.ibm.com

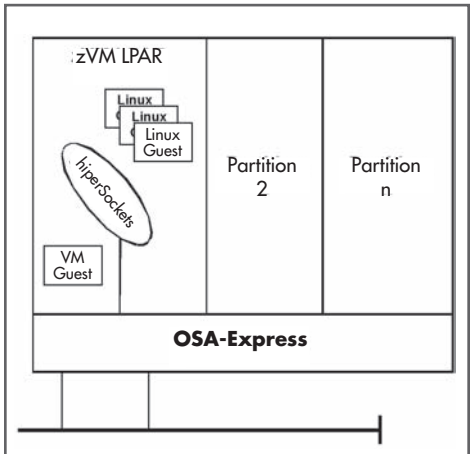

FIG 2: PARTITIONED SYSTEM

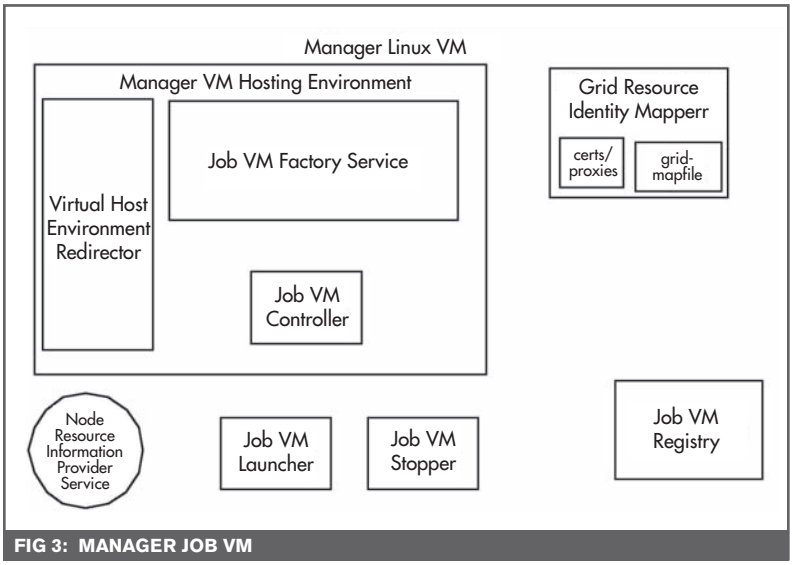

FIG 3: MANAGER JOB VM

Controller searches the Job VM Registry for a corresponding subject. The Job VM Registry is a database where information about active Job VMs on the node is maintained. It replaces the jobmanPortMapping file of Globus. If an entry is found in the Job VM Registry, the target URL is constructed and returned to the Redirector. If an entry does not exist in the Job VM registry, an idle Job VM is selected and the launch module is used to allocate the required resources to it and start it up. The target URL is returned to the Redirector after ensuring the Job VM is running. The active Job VM Registry is updated with this entry.

The Job VM Launcher is a shell script that implements a sequence of operations described later. It is invoked by the Job VM Controller to allocate resources to start a new Job VM. The Job VM Stopper is a shell script that implements a sequence of operations described later. It is invoked by the Job VM Controller to stop a Job VM and deallocate resources assigned to it. The Node Resource Information Provider Service is a specialized notification service providing raw data about a resource characteristic, the file system, the host system, etc.

Most of the components in a Job VM are based on the corresponding GT3 components and services as shown in Figure 4.

The Job VM Hosting Environment is the equivalent of the Globus User Hosting Environment with the modification that it runs differently in a Linux VM than in a Manager VM Hosting Environment.

The Managed Job Factory Service (MJFS) is the same as the Globus MJFS; it is responsible for instantiating a new Managed Job Service (MJS) when it receives a CreateService request. The MJFS stays up for the life of the Job VM. The MJS is the same as the Globus MJS; it is an OGSI service that, when given a job request specification can submit a job to a local scheduler, monitor its status, and send notification. The MJS will start two File Streaming Factory Services (FSFS), one for the job's standard output (stdout) and one for the job's standard error (stderr). The MJS starts the initial set of File Stream Service (FSS) instances as specified in the job specification. The FSFSs Grid Service Handles (GSH) are available in the SDE (service data element) of the MJS, which will enable the client to start additional FSS instances of stdout/err or terminate existing FSS instances. The FSS is an OGSI service that when given a destination URL will stream from the local file the factory that was created to stream (stdout or stderr) to the destination URL. The VM Resource Information Provider Service is a specialized notification service providing raw data about the VMs file system, host system, etc.

Operations

The overall sequence of operations is as follows: the Manager VM exposes the capabilities and characteristics of the resources allocated to the LPAR and the current state of the Job VMs running within the LPAR. When a new job request is received, the Manager VM allocates the necessary resources to a predefined Linux VM and starts the VM. It passes the job request to this new Linux VM and returns to the grid management service a handle to communicate with this VM. The Linux VM executes the tasks. During this period, the grid management service can query for the status of the job and can retrieve the results of the job when ready. Upon completion of the job, the Manager VM shuts down the Job VM, cleans up the used resources, and reclaims those resources.

Figure 5 further details these interactions and relates them to the grid user actions. The user submits a job request to the grid management service; the job request contains information about the program to be executed, the data that the program must operate on, and the resources needed to execute the program. The grid management service queries the Manager VMs running on the different nodes for the availability of the required resources. Each Manager VM responds with the available resources on the corresponding node that it manages. The Grid Portal selects the appropriate node and submits the job request to the corresponding Manager VM. The Manager VM activates one of the preconfigured Linux VMs on the node and allocates the necessary resources to

CTIA WIRELESS 2005

A Division of CTIA–The Wireless Association™

The Most Important Technology Event of the Year!

March 14-16, 2005 Ernest N. Morial Convention Center New Orleans, LA **www.ctia.org**

REGISTER BEFORE JANUARY 31, 2005 and Save up to $345!

CTIA's 12th Annual Wireless Gala

The wireless industry's premiere networking event is always the hottest ticket in town!

For sponsorship information contact Bonnie Knight at bknight@ctia.org.

Keynotes

Stay tuned for more exciting Keynote announcements...

Daniel A. Carp
Chairman and Chief Executive Officer
Eastman Kodak Company

Jeong Han Kim
President
Samsung Telecommunications America

Steve Largent
President/CEO
CTIA-The Wireless Association

CTIA WIRELESS 2005

is the most important global event where standards bodies converge to set the agenda for the future of the wireless industry. Network with the largest gathering of wireless engineers and technologists to find the tools you need to help advance the wireless industry.

CTIA WIRELESS 2005 offers numerous opportunities for the wireless engineer and technologist:

- **CTIA Education** – *a track dedicated to wireless technology*
- **Tower Summit & Trade Show 2005** – *premiere wireless infrastructure event*
- **IEEE Wireless Communications Network Conference (WCNC) 2005** – *the industry's foremost conference for developing wireless standards and engineering*
- **Developers Conferences**
- **Exhibits** – *a 400,000 sq. ft. floor displaying the latest in wireless technology*
- **Wireless News** – *More new wireless product announcements than any other show!*
- **CTIA Wireless Home** – *a 7,000 sq. ft. home that brings wireless technology to life*

Produced by

CTIA
The Wireless Association™

that VM.

The Manager Job then passes the job request details to the Job VM and starts execution of the job. The Manager VM returns a handle to the Job VM to the Grid Portal. The handle enables the Grid Portal to communicate directly with the Job VM. At this point, the job submission process is completed and you are informed of the successful submission of the job.

The Job VM continues the execution of the job. At any time during this process you can query the Grid Portal about the status of the submitted job. The Grid Portal in turn will query the Job VM for the status of the job.

The Job VM notifies the Grid Portal when the job is completed. The Grid Portal retrieves the results and instructs the Manager VM on the same S/390 node to shutdown the Job VM. The Manager VM deactivates the VM, cleans up the allocated resources, and puts them in the available list.

Future extensions of this design will handle more elaborate cases. For example, the user can be replaced by an automated service or program, and instead of the execution of a single job on a single node, the task at hand might require multiple jobs that run simultaneously on multiple nodes.

STARTING A JOB VM

The startup command for a Job VM is as follows:

```
rexec -l vm_userid -p vm_password vm_hostname start
tar-get_userid {-mem mem_size} {-proc proc_num}
```

This command executes a start script, passing it the specified arguments. The first argument specifies the user ID of the target Job VM. The subsequent arguments are optional and are used to indicate that additional resources are needed to process the request. That is, the Manager VM checks the resources defined for the Job VM to ensure that there are sufficient resources to process the request.

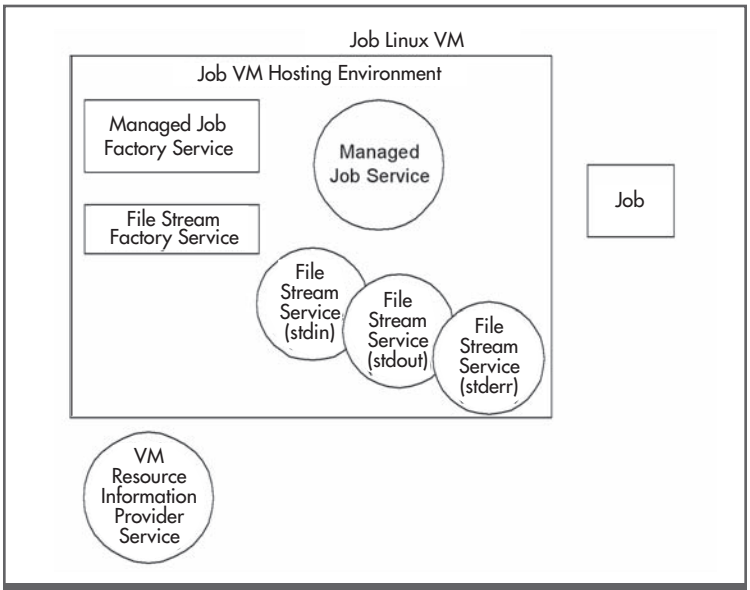

FIG 4: JOB VM ARCHITECTURE

If additional resources are desired, then those resources are requested in this command. For example, -mem specifies the memory size to be allocated, and –proc specifies the number of virtual processors to be allocated.

The start script autologs the specified user ID and issues the appropriate commands to add resources, if needed. Then it IPLs (initial program load) the Job VM. If indicated in the rexec command that resources are needed, the start script issues the appropriate commands to add those resources to the Job VM. For example, if virtual storage is to be added to the Job VM, then a DIRMAINT command with a storage operand, such as DIRM FOR userid STORAGE 1G, is issued. As another example, if the virtual machines needed the maximum virtual storage size available, then a DIRMAINT command with a max store operation, such as DIRM FOR userid MAXSTOR 2048M, is issued.

Should a virtual processor be needed, a DIRMAINT command with a CPU operand, such as DIRM FOR userid CPU cpuaddr, is provided. If a filesystem space is added, issue a DIRMAINT command with an AMDISK operand, such as DIRM FOR userid AMDISK vaddr xxx. In this case, a RACF command is also used to define the disk to RACF. Such a command includes, for instance, RAC REDFINE VMMDISK userid.vaddr OWNER(userid).

In addition to adding the resources to the Job VM, the Job VM is IPL-ed. In one example, this includes reading a name file that is kept for the Job VM instance, autologging the Job VM instance based on the information, and booting up any disks relating to that instance. This completes the startup of the Job VM.

Starting the Job VM also starts the Linux instance configured in the VM. The Linux instance is itself configured to start the Globus container with a predefined set of grid services that includes the modified LMJFS. At this point, the Job VM is operational and ready to accept jobs.

STOPPING A JOB VM

The shutdown command is as follows:

```
rexec vm_userid -p vm_password vm_hostname shutdown
tar-get_userid
```

In response to receiving the command, a shutdown command is issued to the Job VM. Any additional resources allocated to the Job VM are returned. For example, this can be accomplished by issuing the appropriate DIRMAINT/RACF commands, which depend on the type of resources to be returned. If the resource to be returned is virtual storage, then the DIRM FOR userid STORAGE 512M command is issued to return the virtual storage level to its original amount. Similarly, if virtual processors are to be returned, then a DIRM FOR userid CPU cpuaddr DELETE command is issued to delete a virtual processor. As for the delete the filesystem space, a DIRM FOR userid DMDISK vaddr command is issued.

Additionally, cleanup of the Job VM is performed. The

cleanup includes removing old files and placing the Job VM back to its original image.

Benefits

Our design addresses the need for a grid solution for the hosted environment using technologies available today. It also addresses the strategic need to converge autonomic computing, OGSA, and systems design to enhance performance with the optimization and efficiencies that help data centers reduce costs through a pooled set of resources and the ability to host capacity on-demand for a select set of grid applications. The design is set to link mainframe servers from across the company (or company to company) into one, highly utilized grid to help cut application runtime and bring the results even faster. Our design also allows reaching out and grabbing mainframe white space from around the world when needed and giving it back when it is not needed. Finally, it provides better isolation between grid jobs, thus enhancing privacy and security in a grid environment.

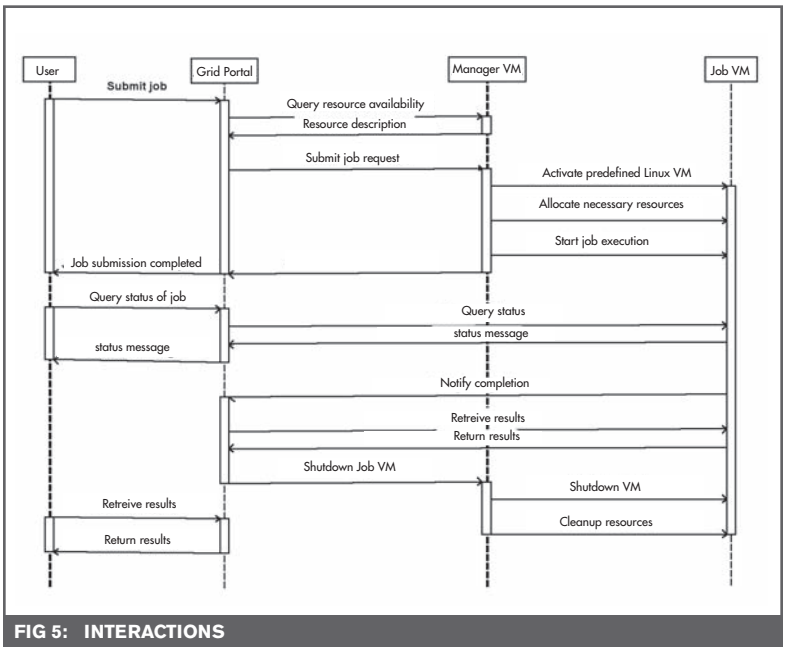

FIG 5: INTERACTIONS

INCREASE IN UTILIZATION

Our design provides interoperability among the different processor families. It uses unused or excessive mainframe computing resources and operates independently while the system is in use. While mainframe Sysplex works within the same family of processors and operating systems, our works with all families and operating systems. Thus it goes beyond the capabilities that Sysplex and Geoplex provide in clustering S/390 systems belonging to only one family of processors.

A vast quantity of computing power is wasted due to the underutilization of resources. In general, planning and sizing for computing requirements are based on peak demand and statistically the actual utilization is in the

order of 60%. Harnessing the unutilized compute power will provide immediate economic benefits to any organization that has a large installed base of servers. White space is defined as the unutilized capacity or cycles on S/390 or z/Series machines. Basically, users on a VM or MVS system only use part of the maximum capacity of the systems so there is room for more workload. In our proposed design, white space is utilized by adding Linux virtual machines.

ISOLATION BETWEEN JOBS

In general, the isolation between grid tasks is based on whatever isolation mechanism an operating system on those platforms can provide. Usually grid tasks are run as separate processes within the operating system, thus sharing resources controlled by the OS. This situation might result in intentional or accidental exposure or corruption of the data of one task by the other task.

Our design exploits the capability of running multiple concurrent virtual machines on a LPAR of a mainframe node to execute each task in an individual virtual machine.

In our design, the isolation mechanisms of GT3 have been replaced by a different model. In GT3, the isolation is based on running a separate Java VM for each user that hosts a Local Managed Job Factory Service (LMJFS). In this design, instead of running multiple UHE instances in the same Linux instance, each UHE instance is executed in a separate Linux instance. Consequently, jobs belonging to different users never share resources.

Summary

Our design provides interoperability among the different mainframe processor families. It uses a part of the mainframe's (S/390 and zSeries system family) unused or excessive computing resources via the logical partitioning (LPAR), and operates independently while the system is in use.

The design is based on porting existing grid computing standard architecture components to the zVM virtual Linux environment, while adding automated setup/configuration features and exposing the resources available on the IBM mainframe nodes through the grid environment's information indexing infrastructure.

The porting of the Globus Toolkit to the zLinux environment has been completed. The automated configuration of each Manager Linux VM through a registration portal is accomplished. A number of business and engineering applications have been tested on the prototype system. The mechanisms to dynamically start and allocate resources to the Job VM have been designed, implemented, and tested. The next step is to integrate these mechanisms with the modified Grid Resource Management services to complete the implementation of the design.

This mainframe grid infrastructure provides IBM increased flexibility and enables the user to tap some of the data center white space computing power by interconnecting mainframes in Poughkeepsie, New York; Boulder, Colorado; London; Tokyo; and Sydney, Australia.

"The design we offer provides better isolation between tasks – something not available in grids based on other platforms"

References

- I. Foster, C. Kesselman, S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *International Journal Supercomputer Applications*, 15(3), 2001.
- *IBM Grid Toolbox version 3 for Multiplatforms*: www-1.ibm.com/grid/solutions/grid_toolbox.shtml
- *Globus Toolkit Version 3*: www-unix.globus.org/toolkit/download.html
- IBM, *z/Architecture Principles of Operation PDF files*, SA22-7832-00 December, 2000: www-1.ibm.com/servers/eserver/zseries/zos/bkserv/zswpdf/zarchpops.html
- I. Foster, C. Kesselman, J. Nick, S. Tuecke, "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," Open Grid Service Infrastructure WG, Global Grid Forum, June 22, 2002, www.globus.org/research/papers/ogsa.pdf
- I. Foster, C. Kesselman, J. Nick, S. Tuecke, "Grid Services for Distributed System Integration," IEEE *Computer Magazine*, 35(6), 2002.
- GT3 GRAM Architecture: www-unix.globus.org/developer/gram-architecture.html
- IBM *zVM Directory Maintenance Facility Function Level 410 Command*, Version 4, Release 3.0, SC24-60025-03, October 2002, www.vm.ibm.com/pubs/pdf/hcsl4a20.pdf
- IBM *Program Directory for Resource Access Control Facility (RACF) Feature for z/VM*, Version 1 Release 10.0, May 2002, www-1.ibm.com/servers/eserver/ zseries/zos/racf/pdf/RACFZVM_430.PDF

See what everyone is talking about… ***The Application Server Shoot-Out!***

International Web Services Conference & Expo

Keynotes

Matt Ackley
eBay

Ari Bixhorn
Microsoft

Mike Milinkovich
Eclipse.org

web services conference & expo EDGE

WEB SERVICES EDGE 2005 EAST

Register by FEBRUARY 11, 2005 SAVE Up To $200

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

New for 2005…

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

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

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

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

PRODUCED BY: SYS-CON MEDIA
PRESENTED BY: SYS-CON EVENTS

DIAMOND SPONSOR: Microsoft
GOLD SPONSORS: Novell, SAP
ASSOCIATION SPONSORS: Integration Consortium, OASIS, ineta, eclipse
CORPORATE SPONSOR: Java Community Process
APP.-SERVER SHOOTOUT FACILITATOR: burton GROUP
MEDIA SPONSORS: Information Storage+Security Journal, JDJ, WebServices Journal, .NET Journal, queue, CoDe, XML Journal, wldj, asp.netPRO, WebSphere Journal, ColdFusion Developer's Journal, SD Times, MX developer's journal

SPECIAL INSERT: WEB SERVICES EDGE 2005
WWW.SYS-CON.COM/EDGE
FEBRUARY 15-17, 2005
JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

web services conference EDGE &expo

Web Services Edge 2005 East

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

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

Who Should Attend

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

Features & Attractions

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

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

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

SPECIAL INSERT: WEB SERVICES EDGE 2005
WWW.SYS-CON.COM/EDGE
FEBRUARY 15-17, 2005
JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

ices Conference & Expo

Keynote Speakers

Tuesday, February 15, 11 a.m.

Matt Ackley
Senior Director, eBay Developers Program

Web Services for eCommerce

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

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

Wednesday, February 16, 11 a.m.

Ari Bixhorn
Director, Web Services Strategies, Microsoft Corporation

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

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

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

Thursday, February 17, 11 a.m.

Mike Milinkovich
Eclipse.org

An Open Development Platform for Web Services

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

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

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

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

SPECIAL INSERT: WEB SERVICES EDGE 2005
WWW.SYS-CON.COM/EDGE
FEBRUARY 15-17, 2005
JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

DAY 1 FEBRUARY 15

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

DAY 2 FEBRUARY 16

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

DAY 3 FEBRUARY 17

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

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

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

Novell

FREE Web Services Security Tutorial Presented by Novell

Using a Web Services Framework to Build SOA Applications

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

FREE Tutorial Presented by Java Community Process

Developing Web Services Using Java Technologies

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

Java Community Process

Free Tutorials with VIP Preregistration ONLY!

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

Microsoft

FREE
.NET Tutorial Presented by Microsoft

The Next Generation of Visual Studio

Thom Robbins

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

Free Tutorials with VIP Preregistration ONLY!

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

critical sites

FREE
.NET Tutorial Presented by Critical Sites

Security, The New Reality

Patrick Hynds

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

Free Tutorials with VIP Preregistration ONLY!

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

SPECIAL INSERT: WEB SERVICES EDGE 2005
WWW.SYS-CON.COM/EDGE
FEBRUARY 15-17, 2005
JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

JAVA TRACK

J-1 What's New In JDO 2.0

Patrick Linskey, SolarMetric

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

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

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

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

Yakov Fain, Smart Data Processing, Inc.

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

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

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

J-3 The ROI of a Java-Rich Client

Coach Wei, Nexaweb Technologies

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

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

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

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

Sudhir Bhojwani, Sushil Shukla, & Sudhrity Mondal, BEA

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

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

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

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

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

J-5 J2ME and Eclipse

Michael Van Meekeren, IBM

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

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

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

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

Pierre Fricke, D. H. Brown Associates

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

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

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

J-7 Migrating Enterprise Applications Between J2EE Application Servers

Ajit Sagar, Infosys Technologies, Ltd

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

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

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

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

Dr. Michael Juntao Yuan, University of Texas

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

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

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

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

Kieran Taylor, Akamai

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

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

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

J-10 Java Web Services Programming Tips & Tricks

André Tost, IBM

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

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

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

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

Onno Kluyt, Sr. Director & Chair, JCP Program, Sun Microsystems

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

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

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

Attend a FREE* Tutorial

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

Aaron Williams

Developing Web Services Using Java Technologies

Aaron Williams, Manager, Java Community Process Program Office, Sun Microsystems

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

Sun microsystems

* Free Tutorials with VIP Preregistration ONLY!

SPECIAL INSERT: WEB SERVICES EDGE 2005
WWW.SYS-CON.COM/EDGE
FEBRUARY 15-17, 2005
JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

International Web Services Conference & Expo

Register by FEBRUARY 11, 2005 SAVE Up To $200

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

CONFERENCE & EXPO: February 15 – 17, 2005

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

THREE WAYS TO REGISTER FOR CONFERENCE

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

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

Please complete sections 1, 2, 3 and 4

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

First Name _____ Last Name _____

Title _____

Company _____

Street _____

Mail Stop _____

City _____

State _____ Zip _____ Country _____

Phone _____

Fax _____ E-Mail _____

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

☐ Check or Money Order Enclosed (Registration confirmed upon receipt of payment)

Check # _____ Amount of Check $ _____

Charge my ☐ Visa ☐ MasterCard ☐ American Express ☐ Discover

Name on card _____

Card # _____ Exp. Date _____

Signature _____

Billing Address (if different from mailing address)

_____

3 PLEASE INDICATE YOUR CONFERENCE CHOICE:

Total Registration fee $_____

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

SPECIAL GROUP RATES

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

4

A. Your Job Title

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

B. Business/Industry

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

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

	Location	Company
10,000 or more	01 ☐	01 ☐
5,000 - 9,999	02 ☐	02 ☐
1,000 - 4,999	03 ☐	03 ☐
500 - 999	04 ☐	04 ☐
100-499	05 ☐	05 ☐
100 or less	06 ☐	06 ☐

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

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

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

E. What is your company's gross annual revenue?
F. Do you recommend, specify, evaluate, approve or purchase mainframe products or services for your organization?
01 ☐ Yes 02 ☐ No

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

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

_____

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

CANCELLATIONS, SUBSTITUTIONS, REFUNDS

Fax written request to SYS-CON Registration 201-782-9601. Requests for refunds received prior to January 7, 2005 will be honored, less a 10% handling charge; requests received after January 7, 2005, and before February 4, 2005, will be honored less a 20% handling charge. No requests for refunds will be honored after February 4, 2005. Requests for substitutions must be made in writing prior to February 11, 2005. No one under 18 is permitted to attend. No warranties are made regarding the content of sessions or materials.

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

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

SPECIAL INSERT: WEB SERVICES EDGE 2005

WWW.SYS-CON.COM/EDGE

FEBRUARY 15-17, 2005

JOHN B. HYNES CONVENTION CENTER • BOSTON, MA

Legacy Is the New Thin

BY JOE WINCHESTER

Paul Simon sings, "Every generation throws a hero up the pop charts." Each person who attempts to conquer the highly fickle music or fashion market frequently does so by merely rehashing old ideas. Trends are repeated and what was once passé becomes fashionable again, as the definitions of old and new are blurred by memory and packaging.

Software is no different. Each new wave is set in motion by realizing the failings of the current technology and then trying to solve those failings. Rather than progressing in a forward direction, however, what often occurs is the very forces that drove the latest technology are also found to be lacking, and so there is a return to the old ideas. In some ways this is no different from politics where, if an incumbent is thrown out on the weakness of his/her current form, the opposition merely returns and attempts the ideas that failed originally. Likewise with technology (although with each reelection time has moved on), hardware is quicker, networks are faster, and problems that let the previous incarnation of the solution down no longer have the same relevance.

The largest, most obvious example of this pendulum swing is the Internet. In the 1970s, mainframes ruled the world; computing was done on a large box with users working on attached terminals; scalability was achieved through the shared loading of programs, pages were cached for each job's session, scheduling was done by queuing time-sliced work, and prestarted resource pools helped with performance. Two evolutionary predators threatened the beast: (1) Moore's Law allowed faster computing power to be provided at lower costs, and (2) the demand for the graphical user interface. PCs ruled the GUI. However, on their own they couldn't provide high-end computing for multiuser environments, thus the client/server model was born – a topology in which the front-end PC had the computing resources to power the GUI while the back end was the workhorse for the application.

A problem arose as to how big a horse the back end would be, with many companies having different vested interests in each endpoint of the equation. Wherever the logic was, so would go the processing dollars. To bolster the front-heavy version the terminal was called "dumb" and, with the subtle inference that it would one day be replaced entirely, the server was nicknamed "legacy." Initial implementations attempted to front-load all of the logic with the back end merely answering SQL requests. The voices of those who had the experience and knowledge to know this solution would never scale were somehow never heard as the emperor paraded his ideas.

The lure of the GUI wasn't enough to sustain unsafely practiced client/server though. All that glitters isn't necessarily gold, and scalability and distribution became the new issues. To move forward from here the software pendulum swung back toward the server. HTML and HTTP were in the right place at the right time and replaced the dumb terminals' display formats. Green was replaced by color, and white became the new black.

The topology of the Internet is remarkably similar to that of the mainframe; however, the previous technology swing had been spent attempting to market them to the Stone Age, so new buzzwords were required. Client/server was named "fat" and what was once called "dumb" was labeled "thin," and the fashion metaphors became more than obvious.

When the client/server problems of distribution were solved, however, the issue of how to scale large applications across the Internet still remained. For this, the transaction servers of yore were resurrected and, among the technologies to rise from the ashes, J2EE was born. Servlets are to HTML what mainframe programs are to their display buffers; a session bean with its state is analogous to page-shared memory in a polled buffer. The mainframe computing architecture was fashionable once more. In a world where networking was prevalent and, ironically, client machines cheaper and more readily available, it succeeded beyond all expectations.

The Internet isn't a success because it is a great technology; it was in the right place at the right time. Some of the best ideas in history are a re-presentation of old ones, but with the benefit of a different set of environmental variables in which to test them and avoid the problems they initially suffered the first time around. There is no shame in this, and it's interesting to note that the big players in the Web market are those with a proven track record in large-scale server computing.

The lesson to be learned in all of this is that what worked for your parents and boss, both in music and in the cut of their pants, might work for you, too. Technology likewise and, with the benefit of hindsight and previously unavailable, cheaper, and faster resources, problems that plagued past implementations in their previous run up the hit parade can now be dealt with. There is no shame in standing on the shoulders of those who have gone before you and learning from the past. Instead of chasing the latest and greatest specs or industry buzzwords, shouldn't we all look back at what technology lies on the shelf and see if in today's environment it can be reborn and flourish as the new super-fit survivor?

Joe Winchester is a software developer working on WebSphere development tools for IBM in Hursley, UK.
joewinchester@sys-con.com

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.

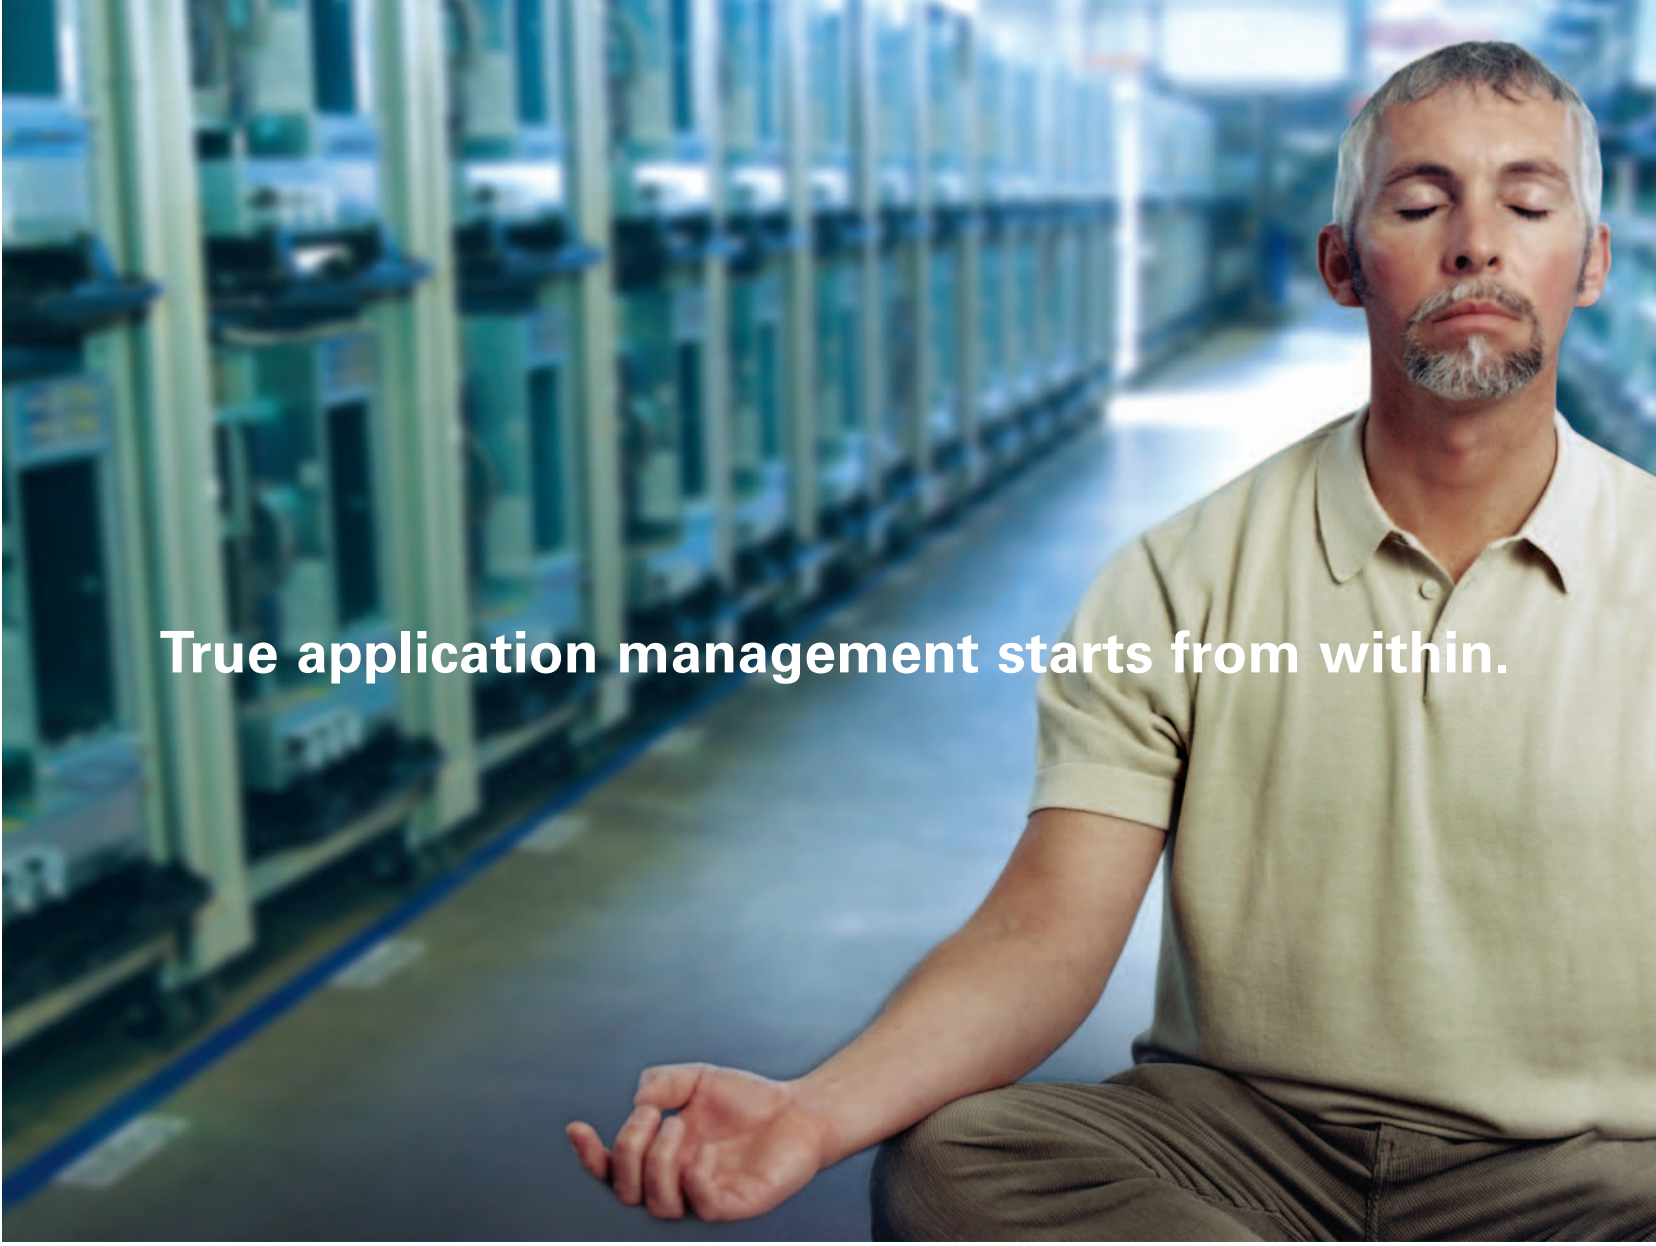

Motive brings an enlightened approach to application management, building management intelligence into the application itself. Only Motive transcends the fragmented, disjointed reality of traditional application management, to deliver the visibility, insight and automation that support and development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