



European Patent Office: Moving to an Open Platform via IBM Technology and Java

An IDC e-business Case Study

THE SUBJECT

With offices in The Hague, Munich, Berlin and Vienna, the European Patent Office (EPO) provides intellectual property protection in 19 European nations as well as other signatory nations. Since 1978, the organization has published more than a million patent applications.

THE GOAL

To migrate the EPO's IT infrastructure from a proprietary client/server platform to an open standards-based, Java-centric three-tiered architecture—enabling the EPO to more easily offer a broad range of Web-based services to its key stakeholders.

THE SOLUTION

The EPO's solution is a suite of Java-based document management applications used by the EPO's internal staff of examiners and the general public, respectively, to search and view any of the organizations 40 million patent-related documents. Developed using IBM VisualAge for Java, this infrastructure-oriented solution relies heavily on IBM WebSphere Application Server running on IBM RS/6000 servers, which are linked to an IBM S/390 server running IBM DB2 at the back end of the solution.

WHY IBM

"We were looking for a partner who could give us the architectural advice and support we needed as we migrated towards Java and platform neutrality. We saw IBM as a partner who, when the going got tough, would stand by us."



Table of Contents

Executive Summary	1
Situation Analysis	2
Background	2
The Need: A Flexible—Yet Industrial Strength—IT Infrastructure	2
Action Plan and Decision Process	3
First Steps	3
Challenges	4
Exhibit: Challenges at Various States of the EPO's e-business Evolution	4
Solution Profile and Implementation Strategy	5
A Brief History of EPO's IT Infrastructure	5
The Project: The Path from OS/2 to Java	5
Exhibit: Key Events in the Development of the EPO's Java-based Solution	5
EPO's Java-based Services In Action	6
Solution Architecture	7
Exhibit: Basic Architecture of the EPO Solution	7
Business Results	8
Exhibit: Overview of EPO's Business Results Achieved	8
Case Epilogue	9

Executive Summary

Innovation Spotlight

The Java support provided by WebSphere Application Server was crucial to achieving the EPO's most important architectural goal: platform neutrality. "WebSphere enabled us to create platform neutrality on the server side as well as on the client side," says David Allin, Director Research & Prototyping at the EPO.

The EPO's e-business solution is comprised of a suite of Java-based document management applications targeted to both internal EPO staff and external users, such as patent lawyers and potential patent applicants. The EPO's **Jviewer** application—used by internal EPO staff to search and examine any of the EPO's 40 million text-based and image-based patent documents—employs a Java application at the client level. The EPO's **File Inspection** and **esp@cenet** applications, which use Java applets at the client level, provide the public with Web-based access to the EPO's patent-related archives, including published patent applications and administrative files related to applications.

The solution, which replaced a client-server based application used by internal staff, was deployed as part of a broader IT initiative in which the EPO adopted a three-tiered architecture, with IBM RS/6000 servers in the middle tier and an IBM S/390 server in the back end. Architected by IBM, the solution stands out because it so radically improved the flexibility of the EPO's IT infrastructure with absolutely no disruption of service to its customers. Given the huge volume of patent applications handled by the EPO, this was no small feat.

European Patent Office Solution at a Glance

- ▶ **Core Functionality** Java-based document management applications used by internal EPO staff and external users, respectively, to search and examine any of the EPO's 40 million text-based and image-based patent documents.
- ▶ **Software** IBM WebSphere Application Server 3.5.2 for AIX, IBM VisualAge for Java, IBM DB2 Universal Database 6.1, IBM CICS
- ▶ **Servers** IBM RS/6000, IBM S/390 Parallel Enterprise Server
- ▶ **Services** IBM designed the architecture and provided key insight in developing the solution's core Java-based applications.
- ▶ **Key Benefits**
 - ▶ EPO expects its Web-based services to save patent applicants approximately 10 million EUR (\$US 9 million) annually.
 - ▶ Web-based services have dramatically improved the quality of service and have substantially shortened the patent application cycle.
 - ▶ Using VisualAge for Java allowed for a faster, more efficient development process through such side benefits as a lower rate of coding errors.
 - ▶ Moving from OS/2 to Java will lower application support costs since the EPO is now only required to maintain a single code base.
 - ▶ Adopting a platform-neutral architecture facilitates data sharing with country-specific patent offices, while a 3-tiered solution allows fast roll-out of new applications.

► Background

The European Patent Office (EPO) is an inter-governmental organization whose principal line of business is the protection of intellectual property or, more specifically, inventions. While its core jurisdiction extends to 20 European nations, the EPO also provides protection within other, non-European countries that have signed agreements with the EPO. Since processing its first patent application in 1978, the organization has published more than a million applications and presently publishes new patents at a rate of 2,000 per week—and the volume is growing. At its most basic, the EPO's mission is to compare any new invention for which a patent is applied against all published inventions since 1920, and to either declare it truly “novel”—and worthy of 20 years patent protection—or reject it. With applicants ranging from individual inventors to corporations such as Siemens and Unilever, the outcome of the patent review process can have enormous significance—in some cases determining a company's survival.

Since processing its first patent application in 1978, the EPO has published more than a million applications and presently publishes new patents at a rate of 2,000 per week—and the volume is growing.

While the stakes are high, the data-processing requirements are even higher. To perform its core analytical functions, EPO maintains a vast archive made up of more than 30 million documents containing patent-related product data and another 10 million supporting documents. What's more, this enormous storehouse of documents is constantly growing, with the EPO receiving thousands of pages of patent applications daily, each of which is scanned, converted to text and—after being checked for quality—stored in its IBM S/390. To perform the required document search and examination functions, the EPO maintains a staff of approximately 5,500 in Munich, The Hague, Berlin and Vienna. In 1999, EPO staffers conducted a staggering 125,000 novelty searches. “The EPO obviously has a very clear need for advanced information technology,” notes David Allin, Director Research & Prototyping at the EPO. “In 1998, it became equally clear that we needed to take action to bring our IT infrastructure in line with an emerging set of needs—both inside and outside our organization.”

► The Need: A Flexible—Yet Industrial Strength—IT Infrastructure

In early 1998, faced with ever-growing data processing requirements and an aging IT infrastructure, the EPO's IT planning staff saw the need to establish a firm footing for the future as imperative. The core issue affecting EPO's technology strategy was the proprietary nature of its legacy system, the most important element of which is a client-server system employing a fat OS/2 client running C++ code. The system—the result of more than a decade of development and modification—accesses the EPO's extensive data repository stored on an IBM S/390. While the system continued to serve their internal needs, a number of external forces made the need for change apparent, notes Allin. “Although we built the system for use inside our own office, we were hearing a chorus of requests from the sovereign patent offices in individual European countries that we grant them access to the information in our database,” says Allin. “However, the incompatibility of our systems made that goal all but impossible.”

At the same time outside parties were clamoring for access to its databases, the EPO also recognized the need to lessen its exposure to the OS/2 platform. While OS/2 had performed well for EPO, the prospect of steadily declining vendor support convinced the EPO's IT department that an exit strategy was required. Thus, while the EPO had clearly recognized and accepted the need to change the core elements of its IT infrastructure, it had yet to resolve what exactly the next chapter of its evolution would be.

Action Plan and Decision Process

► First Steps

To better understand their range of infrastructure options, EPO planners met with a number of solutions providers, including IBM, Sun Microsystems, and various consulting organizations. In the course of their discussions, the EPO team narrowed their selection to either building another client/server application using Windows as the operating system, or adopting a platform-neutral, Web-based model. As Allin explains, the Java-centric, platform-neutral strategy advocated by IBM emerged as the clear and early favorite because it allowed them to avoid an “all or nothing” approach. “If we moved to a Windows model, we would have to deploy *all* of the applications before we could use any of their applications,” says Allin. “We were also very nervous about the prospect of having to redevelop the entire application suite. The platform-neutral approach allowed us to migrate our applications component by component—at our own pace—without taking on excessive risk.”

“We were looking for a partner who could give us the architectural advice and support we needed as we migrated toward Java and platform neutrality. We saw IBM as a partner who, when the going got tough, would stand by us.”

— David Allin, Director
Research & Prototyping,
EPO

In addition to advocating platform neutrality and Java, the IBM team proposed that the EPO reorient its application architecture from client server to a three-tiered model. According to Allin, the team embraced the idea of a tiered architecture because it fit perfectly with the EPO's goal of keeping its hands off its mainframe systems. “The fact that a three-tiered architecture effectively provides an isolation layer between the client and back-end means we won't have to change our mainframe systems,” notes Allin. “This gives us room to maneuver that we would not have had if we merely ported the application from one platform to another.”

The EPO selected IBM as its solutions partner in mid-1998, on the strength of its Java expertise, its depth of expertise in planning platform-neutral architectures and the strength of its support. As Allin explains, the complexity and importance of the EPO's IT systems—and the need to integrate with existing mainframe services—pointed to IBM as the clear choice. “We were looking for a partner who could give us the architectural advice and support we needed as we migrated toward Java and platform neutrality,” says Allin. “We saw IBM as a partner who, when the going got tough, would stand by us.”

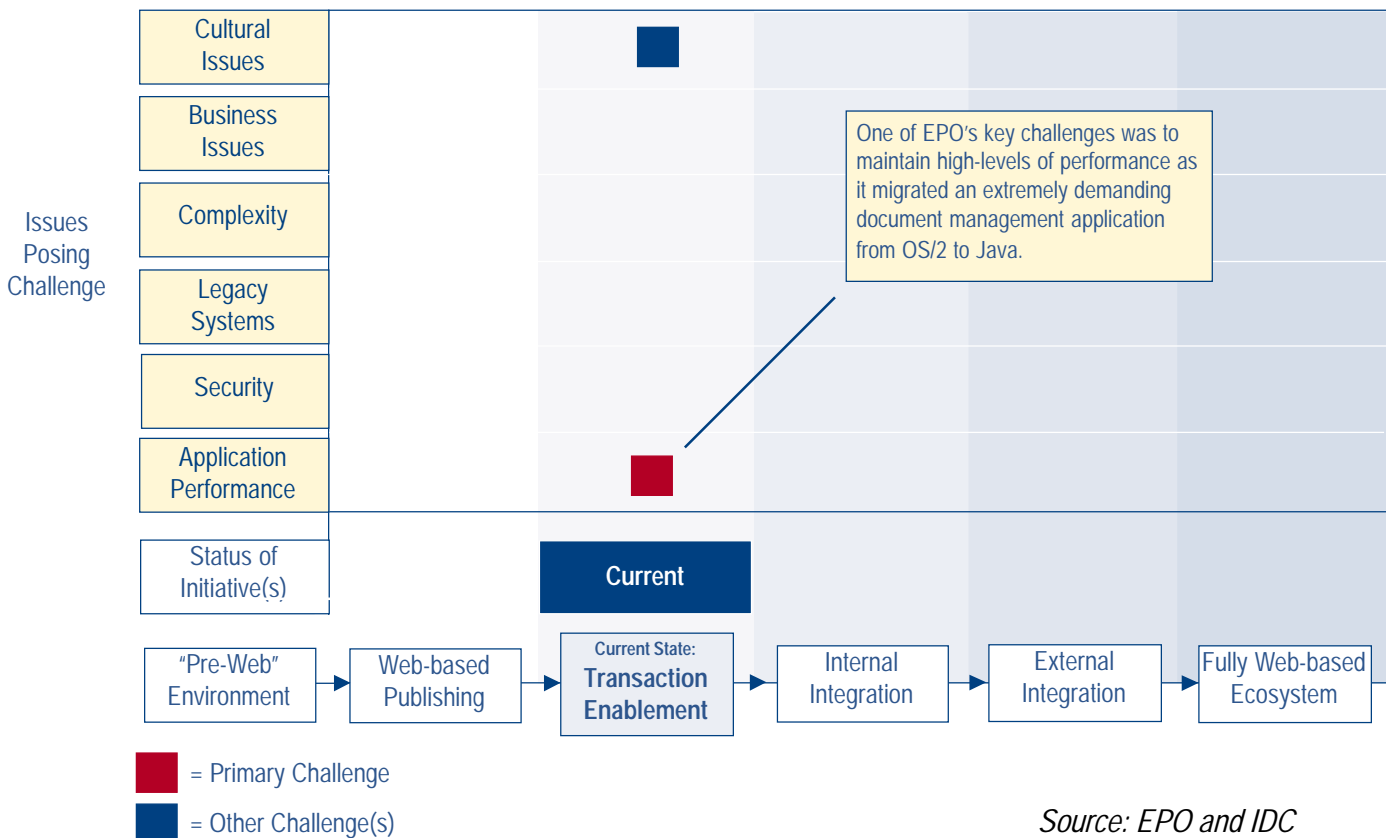
► Challenges

Embarking on its migration from OS/2 to an open, Java-based architecture, the EPO was cognizant of the risks—both technical and cultural. At a technical level, the key challenge was to ensure that the EPO’s demanding set of applications would perform as well or better under the new Java-based architecture. Under one particularly demanding application, EPO examiners:

- perform searches
- download large volumes of compressed image data from the mainframe as a result of these searches
- decompress the images, and
- display them rapidly (0.5 seconds) on the workstation.

As Allin points out, it was imperative that these high levels of performance—the result of years of development in the OS/2 environment—be replicated in Java from square one. “One of our biggest fears was that the users would say ‘why are you fixing it if it isn’t broken?’” says Allin. “Put simply, we had to develop a solution that performed as well or better than what it was replacing.”

Challenges at Various States of the EPO’s e-business Evolution



Solution Profile and Implementation Strategy

► A Brief History of EPO's IT Infrastructure

The EPO's three-tiered IT infrastructure rests atop an IBM S/390 server running IBM DB2, full-text searchable proprietary databases and CA-IDMS which are all used to store patent-related data, and IBM CICS transaction server software, which updates, manages and controls access to the EPO's legacy databases. [Although the EPO formerly ran exclusively Computer Associates' IDMS database management software on its S/390 server, it is currently shifting to DB2 to accommodate its growing need for superior performance, capacity and scalability.] The EPO's older OS/2-based client/server application was developed by a consortium of European solutions providers, with IBM serving as its principal core technology provider, notes Robert Dröge, Director of Application Development at the EPO. "We have a very demanding IT infrastructure," says Dröge. "The strength of IBM technology and the expertise of IBM's professionals were crucial to its construction."

► The Project: The Path from OS/2 to Java

By design, the EPO's OS/2-to-Java initiative has proceeded at an evolutionary pace, reflecting the organization's aim of incremental deployment. After conducting a study in the first half of 1998 designed to ascertain the feasibility of migrating from OS/2 to Java, the EPO gave IBM the green light to design the architecture of the solution. After completing the architecture planning stage

Key Events in the Development of the EPO's Java-based Solution

	3Q98	4Q98	1Q99-4Q99	2Q00-3Q00	4Q2001
After completing a feasibility study, IBM begins laying the architectural foundation for EPO's move from OS/2 to Java.	■				
IBM completes the design of EPO's new three-tiered Java-based architecture.		■			
Development of the Jviewer (1Q99) and File Examination (4Q99) applications begun.			■		
Jviewer and File Examination applications completed (2Q00); development of esp@cenet® application begun (3Q00).				■	
Production version of esp@cenet application to be completed.					■

Source: EPO and IDC

at the end of 1998, IBM provided additional advice related to application development practices that ultimately proved to be very valuable. According to Dröge, IBM's guidance that the EPO follow an incremental, "proof-of-concept" approach to rolling out Java-enabled applications did much to dispel concerns about Java's ability to supplant OS/2 for mission-critical applications. "While the value of this approach is evident now, it was not obvious at the time," says Dröge. "We saw IBM's ability to demonstrate its value as a clear reflection of their knowledge of the architecture all the way down to mainframe level."

In early 1999, the EPO's solutions partners (a consortium of professional service organizations) began the development process, with the initial stages being proof of concept testing for key architecture elements. EPO internal staff served to manage the project. The team's first deliverable—a document management application known as Jviewer that employs a Java application as a client—is used by internal EPO staff to conduct novelty searches among 40 million text and image documents. Development of Jviewer was begun in 1Q99 and completed in 2Q00.

The usage rate of the EPO's Jviewer application—8.5 million text and 10 million image files viewed per month—underscores the importance of maintaining high levels of performance in a Java environment.

From 4Q99 to 2Q00, the team developed EPO's File Inspection application, using IBM VisualAge for Java as the development tool. File Inspection provides external users - such as patent lawyers acting as an agent of a patent applicant—the ability to view any of 550,000 administrative dossiers (containing 13 million documents) related to specific patent applications through a Web browser. These records include patent applications, search reports and all related correspondence between the EPO and applicants. [File Inspection is commonly used by patent lawyers that conduct research and, based on their findings, advise their clients whether or not to file a patent application.] Upon completing Jviewer and File Inspection, the EPO then began a third Java-based initiative known as the rebuild of esp@cenet level II, in 3Q00. The rebuilt application, which essentially provides Web-based access to the same base of document as the Jviewer application, is expected to be in full production later in the year. [Since Jviewer is an application, the EPO has the opportunity to integrate it with the other examiner tools given its functionality and user interaction. By comparison, esp@cenet is a browser-based application whose somewhat more limited performance and functionality characteristics render it inadequate for internal users.]

► EPO's Java-based Services In Action

The Jviewer application is now used by 2,000 internal EPO staff, who view 8.5 million text and 10 million image files per month. According to Dröge, this large number of page views underscores how important it was to maintain the high levels of performance as the application was ported from OS/2 to Java. "With a page flipping rate of a sub-second per page, Jviewer provides our examiners with the superior performance they need to perform their tasks in as efficient manner as possible," says Dröge. The EPO's patent approval process is set in motion by the receipt of a patent application. Upon receipt of an application, EPO patent examiners will then initiate a search using Jviewer among the EPO's millions of published documents, with the ultimate goal of determining if the proposed invention is "novel." The search procedure is

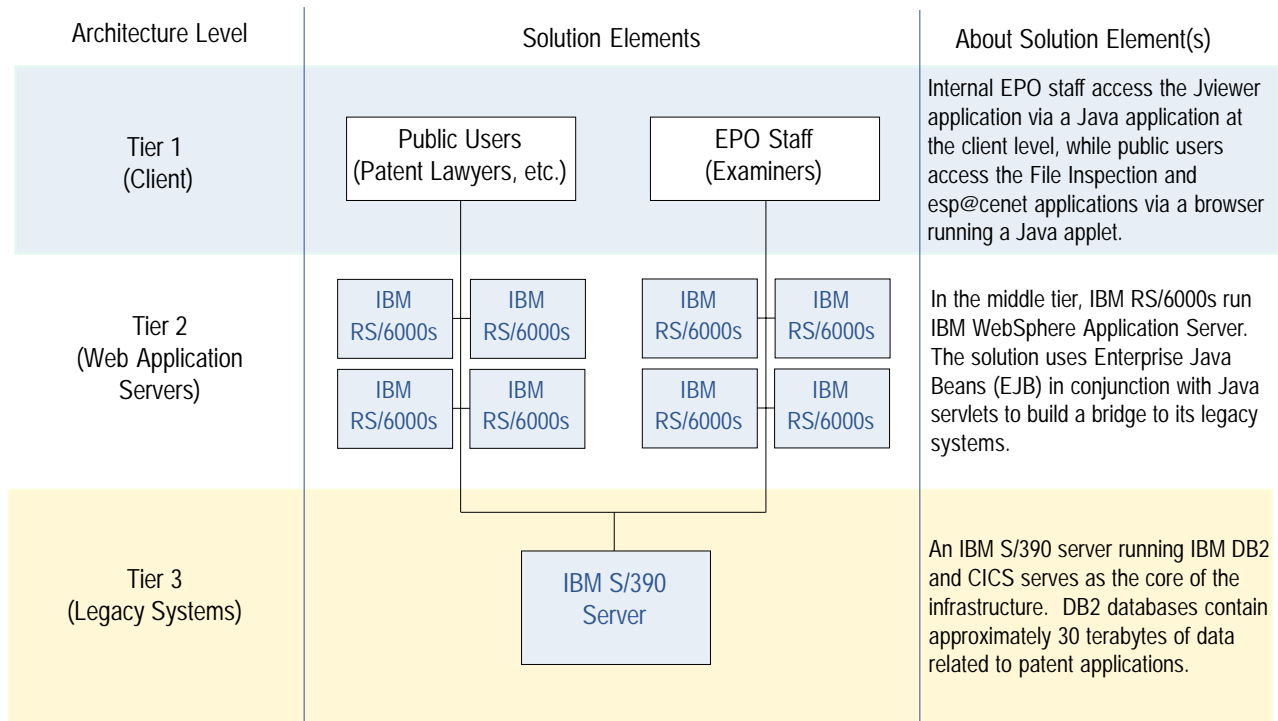
completed with the delivery of a “Search Report” together with copies of the cited documents. Based on the Search Report the applicant can decide whether to proceed with the substantive examination which may eventually lead to granting a patent.

► **Solution Architecture**

The EPO has successfully made the transition from a client server architecture to a three-tiered architecture, with (as discussed) an IBM S/390 server running IBM DB2 and CICS serving as the core of the infrastructure. The main components of the middle tier of the solution are nine IBM RS/6000 servers running IBM WebSphere Application Server—the latter providing the critical middleware link needed to enable Web-based access to EPO’s internal records. At the top tier of EPO’s architecture, the Jviewer client (used by internal examiners) is a Java application, while the File Inspection application is a Java applet running in the browser.

According to Dröge, the EPO leveraged WebSphere’s key functional strengths in deploying Java in the middle tier of its solution. “WebSphere provides the servlet engine that enables us to efficiently handle a high volume of client requests, and to construct the responses that have to be returned,” says Dröge. “Moreover, the use of JSPs (Java Server Pages) enables us to separate the display logic from the business logic implemented in the servlets, giving us a flexible solution for presenting the highly dynamic content of esp@cenet.” WebSphere also provides the Enterprise Java Beans (EJB) container in which to deploy the session beans that are used to facilitate access to the data stored

Basic Architecture of the EPO Solution



Source: EPO and IDC

in the S/390 server. “These beans are application independent,” adds Dröge, “and thus can potentially be re-used elsewhere in the organization.”

Business Results

The EPO’s recent initiatives have produced major business results across a wide range of areas. For the EPO’s external “customers,” the benefits include an enormous increase in convenience and an accompanying decrease in the costs associated with the patenting process. For example, the EPO expects its File Inspection service, which allows Web-based access to EPO data, to save patent applicants and their agents approximately 10 million EUR (\$US 9 million) a year in time and travel costs by eliminating the need to physically travel to the EPO’s offices to research paper-based files. On a more general level, the EPO’s systems have begun to dramatically improve the quality of service delivered to the EPO’s customers, and to substantially shorten the time required to process patent applications.

While a boon to customers, the EPO’s Java-based initiatives have already begun to deliver an equally rich vein of benefits internally. According to Allin, the Java support provided by WebSphere Application Server was crucial to

Overview of EPO’s Business Results Achieved

Business Process Area	Nature of Benefit	Description or Metric
Patent Application (Customers)	Cost Savings	EPO expects its Web-based services to save patent applicants approximately 10 million EUR (\$US 9 million) annually.
Customer Service	Improved quality of service	Web-based services have dramatically improved the quality of service and have substantially shortened the patent application cycle.
Application Development	Increased Efficiency	Using VisualAge for Java allowed for a faster, more efficient development process through such side benefits as a lower rate of coding errors.
IT Administration	Lower Costs	Moving from OS/2 to Java will lower application support costs since the EPO is now only required to maintain a single code base.
IT Administration	Increased Efficiency	Adopting a platform-neutral architecture facilitates data sharing with country-specific patent offices, while a 3-tiered solution promotes speed and flexibility in rolling out new applications.

Source: EPO and IDC

achieving the EPO's most important architectural goal—platform neutrality. “WebSphere enabled us to create platform neutrality on the server side as well as on the client side,” says Allin. “What’s more, WebSphere’s cross-platform support makes it possible for us to develop a new application on one platform such as Windows NT, and then propagate it to the AIX or OS/390 environments. This has been and will continue to be very valuable.”

On the application development side, Dröge provides similar kudos for VisualAge for Java. “VisualAge was an extremely important and valuable tool in the development process,” says Dröge. “Not only did it speed the development process, but it also resulted in a lower rate of coding errors. Overall, it was a very positive experience.”

On a broader scale, the EPO's migration from OS/2 to a Java-based open, platform-neutral IT infrastructure has delivered benefits that will continue to unfold long after the development process is over. One major benefit of the move to Java will be lower personnel costs related to application support, since the EPO is now only required to maintain a single code base—instead of two or three different code bases. Likewise, by adopting a three-tiered infrastructure model, the EPO is now poised to add functionality more quickly, with less risk of disruption, and without having to make changes to its mainframe systems.

Case Epilogue

“It’s safe to say that without IBM’s assistance, we simply would not be deployed right now, which would be bad news for us and for our customers.”

— David Allin, EPO

As it continues to evolve as an e-business, the EPO will introduce a host of new Java-based services, including transactional applications that will enable the online filing of patent applications. Moreover, the EPO plans to provide expanded access to other European patent offices—a task made easier by its embrace of an open, standards-based architecture. According to Allin, these initiatives underscore the EPO's strong emphasis on using e-business technology to improve the depth and quality of services it provides to its constituents. “We’re now in a much better position to provide superior customer service, and to do so in a more efficient, cost-effective manner,” says Allin. “We fully intend to capitalize on the broader range of opportunities that our new architecture affords us.”

Asked about the importance of IBM's contribution to the overall migration initiative, Allin puts it simply: “Without IBM's involvement, it would've been substantially more difficult, would've taken substantially longer, and without their assistance in optimizing the Java applications, it simply would not have been as good a solution,” says Allin. “It’s safe to say that without IBM's assistance, we simply would not be deployed right now, which would be bad news for us and for our customers.”

06-01

AIX, CICS, DB2, DB2 Universal Database, IBM, the IBM logo, OS/2, RS/6000, S/390, S/390 Parallel Enterprise Server, VisualAge and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

Other company, product, and service names may be trademarks or service marks of others.

Printed in the United States of America on recycled paper containing 10% recovered post-consumer fiber.



G325-1846-00