



IBM Software Group

WebSphere® Enterprise Service Bus V6.1
WebSphere Process Server V6.1
WebSphere Integration Developer V6.1

Portfolio mediation tutorial description



@business on demand.

© 2008 IBM Corporation
Updated May 18, 2008

This presentation examines the scenario used for the portfolio mediation tutorials and describes what you will do in each of the tutorials.

Goals

- Understand the portfolio mediation scenario
 - ▶ Describe the original application before the mediation
 - ▶ Discuss reasons for needing a mediation
 - ▶ Describe the application with the mediation
 - ▶ Look at the flow of the mediation
- Introduce the tutorials
 - ▶ Describe the scope of each of the tutorials
 - ▶ Explain the built-in test data



The goal of the presentation is to first understand the portfolio mediation scenario and then to describe the tutorials that are based on it.

An understanding of the scenario is built by first describing the original application before the introduction of a mediation. The business reasons for adding a mediation are discussed, and the application containing the mediation is examined along with the contents of the mediation itself.

Once you understand the scenario, the tutorials that are built upon the scenario are described.

Finally, the test data is outlined to help you understand the expected behavior when running the application.

Section

Portfolio mediation scenario description



In this section, the portfolio mediation scenario is described.

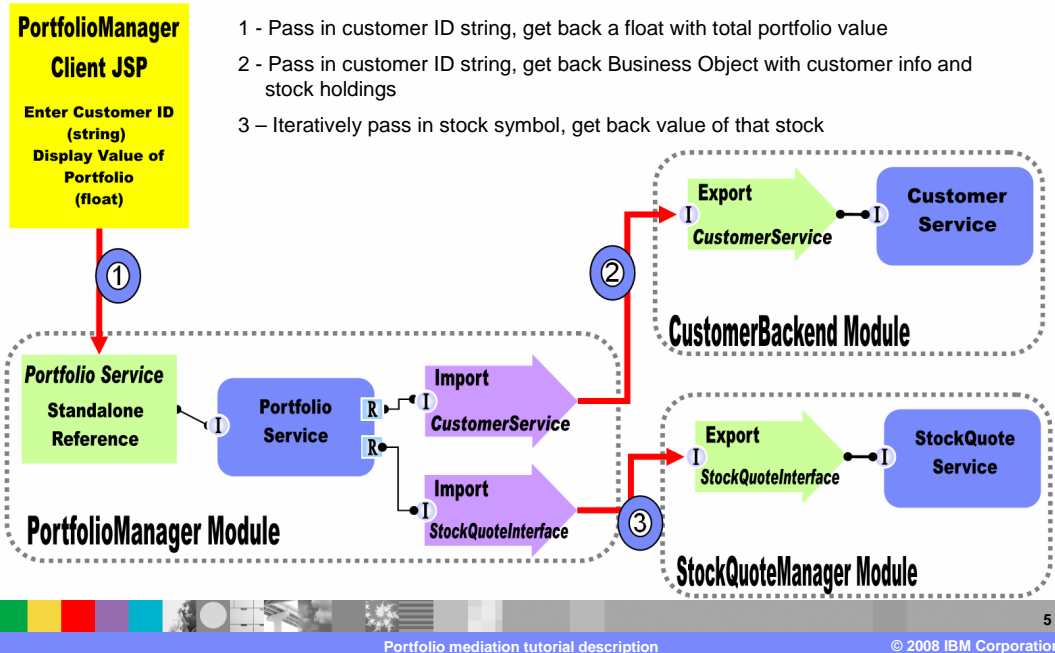
Original SCA application without mediation

- Application computes value of a customers portfolio
- A Java™ Server Page (JSP) provides a browser based interface
 - ▶ Enter customer ID
 - ▶ Total value of portfolio is returned
- Application includes three SCA modules
 - ▶ PortfolioManager module
 - ▶ CustomerBackend module
 - ▶ StockQuoteManager module
- The application flow is illustrated and described on next slide



The application in its original form is used to compute the value of a customer's stock portfolio. It is composed of a JSP which provides a browser based interface where their ID is entered and the total dollar value of their stock portfolio is returned. The application is implemented in three SCA modules, specifically the PortfolioManager module, the CustomerBackend module and the StockQuoteManager module. The following slide shows these modules and the interactions between them.

Original SCA application without mediation



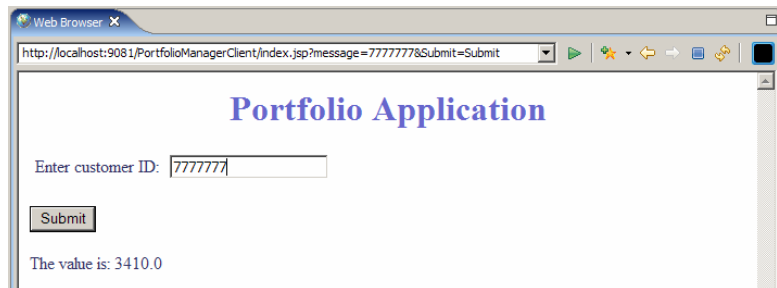
This slide illustrates the original application showing the SCA wiring diagrams for each of the modules along with the basic flow of the application between modules.

Starting in the upper left the yellow box represents the PortfolioManager JSP which provides the browser based user interface where the customer ID is entered and the total portfolio value is displayed. The JSP uses the stand-alone reference in the PortfolioManager module to get a reference to the PortfolioService component and sends the customer ID to the PortfolioService as seen at the number 1.

The PortfolioService then sends the customer ID to the CustomerService component in the CustomerBackend module as seen at number 2. The CustomerService component returns customer information to the PortfolioService, including a list of the customer's stocks and the number of shares held.

The PortfolioService then calls the StockQuoteService component in the StockQuoteManager module as seen at number 3. The StockQuoteService is called once for each stock, with the stock symbol being passed in and the value of one share of that stock being returned. The PortfolioService then computes the total value of the customer's portfolio and returns that to the JSP which displays the value on the browser interface.

User interface to application



- Enter a customer ID
- Click Submit
- Total value of that customers portfolio is returned



This slide shows a screen capture of the browser interface provided by the JSP. As you can see, the interface is very basic. There is a field where the customer ID is entered, a Submit button is used to initiate the processing and the total portfolio value is then returned.

Why is a mediation needed?

- The company has two major goals
 - ▶ Provided enhanced function to the application users
 - ▶ Be able to distribute customer data across multiple new backends



The requirement to add a mediation to this application is driven by a couple of goals the company has for the application. First of all there is a need to provide an enhanced user experience with the capability to present more information to the customer than just the total value of their stock portfolio. The second goal is related more to the underlying infrastructure. The company would like to move away from a single centralized database it now uses and make use of regional databases across which customer data would be distributed.

Why is a mediation needed? (continue)

- Staging implementation of upgrades to meet these goals
 - ▶ Insert mediation module between existing PortfolioService and existing CustomerService backend
 - ▶ Incrementally introduce new CustomerService backends
 - They support an enhanced interface
 - Subset of customer data migrated to each new backend as it comes online
 - Route customer requests to correct backend and handle interface differences



In order to meet these goals, a plan is put in place to stage the changes in such a way that there will be no disruption in application availability.

The first step is to insert a mediation module between the existing PortfolioService in the PortfolioManager module and the CustomerService in the CustomerBackend module. This does not change the application or the infrastructure but rather enables the possibility of mediation logic being added to the flow.

The second step is to provide an enhanced version of CustomerService and introduce that in conjunction with one of the regional databases coming online. The enhanced CustomerService implementation supports a new interface which provides the information required to support the eventual enhancement of the user interface. In addition, a subset of customer data is migrated from the central database to the new regional database. The plan is to stage in the regional databases over time. This will require the mediation to make routing decisions based on the customer ID so that the request is sent to the right backend and will also require the request to be modified to the new CustomerService interface. At this point, the user interface has not changed, but everything is staged as to allow a subset of the customers to be enabled for a new user interface implementation.

Why is a mediation needed? (continue)

- Staging implementation of upgrades to meet these goals (continue)
 - ▶ Introduce a new enhanced PortfolioService and associated mediation
 - ▶ When transition to all new backends complete, drop original PortfolioService and mediation
- Only the first two points are addressed in these tutorials

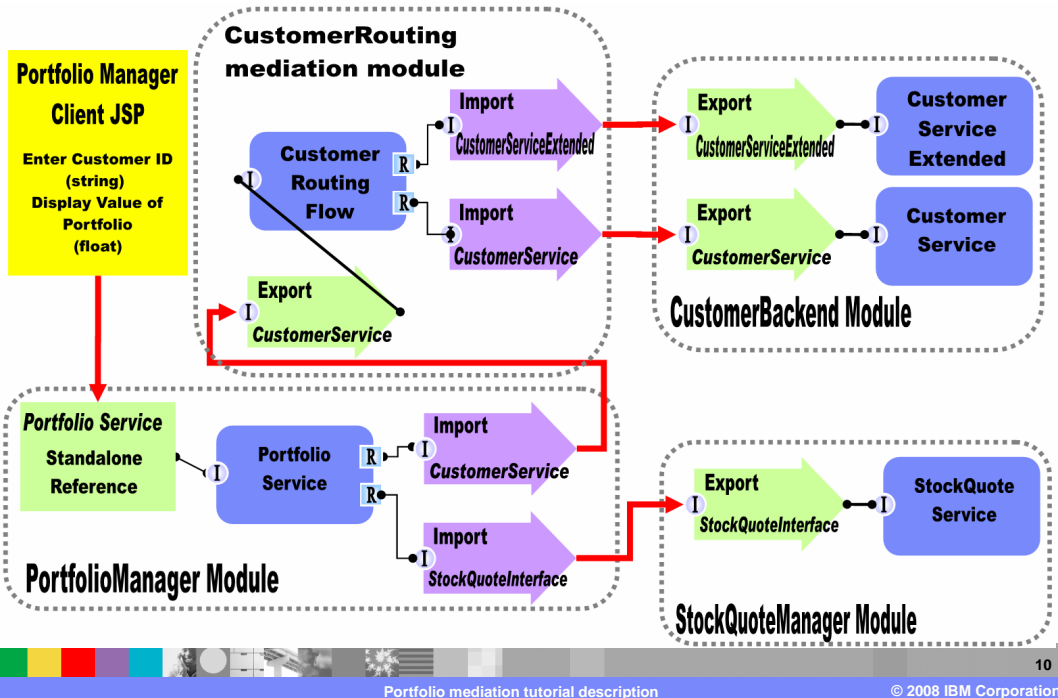


The third step would be to introduce a new PortfolioService and JSP to provide the new user experience along with a new mediation that handles requests for this new application.

The fourth stage would occur once all the customer data has been migrated to regional databases. At this point, the original JSP and PortfolioService can be removed, as can the first mediation.

The tutorials used with this scenario will only address the first two steps in this staging. The following slide shows the structure after the second stage is implemented.

Introduction of the mediation module



Notice that this diagram is quite similar to the original application previously presented. You can see that the interaction between the JSP and the PortfolioService is not changed. Likewise, the interaction between the PortfolioService and the StockQuoteService is also unchanged.

Some differences should be noted:

First, there is now a new CustomerRoutingMediation module in the diagram.

Next, the CustomerService Import in the PortfolioManager module has been modified to connect to the CustomerService Export in the CustomerRoutingMediation module rather than the one in the CustomerBackend module.

Notice that the CustomerBackend module now has a CustomerServiceExtended component and Export in addition to the CustomerService component and Export. This is in support of the new regional database backends.

The function of the CustomerRoutingFlow component in the CustomerRoutingMediation module is to determine the appropriate backend to route each request. Requests to the old centralized backend can be forwarded as is. Requests to the new regional backend need to have the operation modified to meet the CustomerServiceExtended interface.

The next slide examines the logic used in the CustomerRoutingFlow component to accomplish this functionality.

The mediation flow – request flow

- Message logger primitive
 - ▶ Log the request
- Custom mediation primitive
 - ▶ Extract two digit prefix from customer ID in message body
 - These two digits from the customer ID are used to identify backend containing their information
 - ▶ Place two digit prefix into the transient context
- Database lookup primitive
 - ▶ Use two digit prefix from transient context as key for lookup
 - ▶ Lookup returns a backend identifier, places it into the transient context



The CustomerRoutingFlow component starts out with a Message Logger mediation primitive that logs the incoming request.

The next mediation primitive is a Custom Mediation. It is Java™ code that extracts the first two digits of the customer ID from the message body and places them into a field in the transient context. These two digits will be used to determine which backend to use for accessing this customer's data.

A Database Lookup mediation primitive is then used. It takes the two digits extracted from the customer ID and uses them as a key for doing a database lookup that will identify which backend to use, the old centralized database, or a new regional database. The backend is identified with a string value that is placed into the transient context by the database lookup.

The mediation flow – request flow (continue)

- Message filter primitive
 - ▶ Determine routing based on backend identifier
 - ▶ Old backend – go directly to the callout for CustomerService
 - ▶ New backend – go to the XSLT
- XSLT primitive
 - ▶ Transform the CustomerServiceRequestMessage to a CustomerServiceExtendedRequestMessage
 - ▶ Go to the callout for CustomerServiceExtended

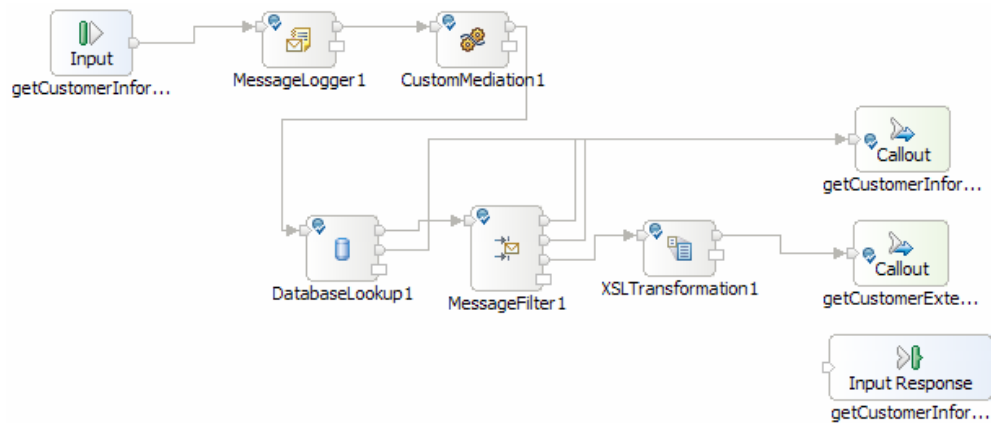


A Message Filter mediation primitive is now used to check the string value indicating the database to use and it selects the appropriate path on which to propagate the message. If the message is to go to the old backend, the path goes directly from the Message Filter to the callout for the CustomerService backend.

In the case where the message is to flow to the new backend, an XSLT mediation primitive is used to transform the message from a CustomerService request to a CustomerServiceExtended request. Once transformed, the message flows from the XSLT primitive to the callout for the CustomerServiceExtended backend.

The following slide will show the mediation flow diagram for this mediation.

The mediation flow – request flow



13

Portfolio mediation tutorial description

© 2008 IBM Corporation


This is a screen capture of the mediation flow described on the previous slide.

On the left side you will see the input node on which the customer service request message comes in and on the right side you see the callouts to the CustomerService and CustomerServiceExtended backends.

The flow starts with the Message Logger going to the Custom Mediation and then to the Database Lookup. On the Database Lookup, if the two digit key is not found, the keyNotFound terminal is used and the flow is taken to the callout for the CustomerService backend. If the key is found the flow goes to the Message Filter. The flow from the Message Filter goes to the callout for the CustomerService backend for two cases. First is when the backend is identified as the old backend and the second is when the filter does not have a matching expression for the backend value. However, when the backend is identified as the new backend, the Message Filter flows to the XSLT to transform the message. The message is then propagated to the callout for the CustomerServiceExtended backend.

The mediation flow – response flow

- CustomerService callout response
 - ▶ Go directly to the CustomerService response node
- CustomerServiceExtended callout response
 - ▶ XSLT transforms from CustomerServiceExtendedResponseMsg to CustomerServiceResponseMsg
 - ▶ Go to CustomerService response node



This is the response flow. When the call went to the CustomerService backend, there is nothing special to do on the return, so the callout response node is wired directly to the input response node. In the case where the call went to the CustomerServiceExtended backend an XSLT mediation primitive is required to transform the response message before returning.

Section

Description of tutorials



Now that a description of the application and mediation has been provided, this section will describe each of the tutorials that have been provided.

Tutorials – Lab1

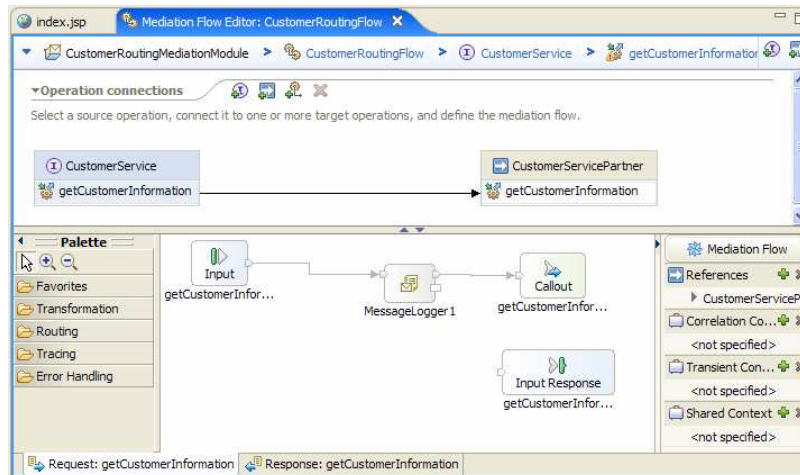
- Start with the original portfolio application
 - ▶ No mediation
 - ▶ Only the old backend with CustomerService interface
- The following are added:
 - ▶ Mediation module
 - ▶ Mediation flow component
 - ▶ A message logger in the flow
- The application is deployed and run
- The message logger database is examined



The first tutorial starts with the original application which uses the old backend that supports the CustomerService interface and there is no mediation. In the tutorial you will create the Mediation Module with corresponding Mediation Flow Component. The mediation flow will only contain a Message Logger that will log the message as it flows through the mediation. This is the equivalent to stage one of the staged approach previously described. The application function has not changed at all but there is now a mediation inserted into the flow.

The application with this mediation will be deployed to the test server in the WebSphere Integration Developer and run by using the browser interface to initiate a request for a customer. The database used by the Message Logger is then examined to see what has been logged.

Tutorials – Lab1



17

Portfolio mediation tutorial description


© 2008 IBM Corporation

The top screen capture shows the assembly diagram for the mediation module. It contains an Export supporting the CustomerService interface wired to the Mediation Flow Component. This in turn is wired to an Import which also supports the CustomerService interface.

The bottom screen capture shows the Mediation Flow Editor with both the operation connections panel and the canvas containing the mediation flow. You can see that the mediation flow is composed only of a Message Logger.

Tutorials – Lab2

- Start with the results of Lab1
- The following are added:
 - ▶ Import to access the new CustomerServiceExtended backend
 - ▶ Expand Mediation flow by adding and configuring
 - Custom Mediation
 - Database Lookup
 - Message Filter
 - XSLT (in both the request flow and the response flow)
- The application is deployed and run



18

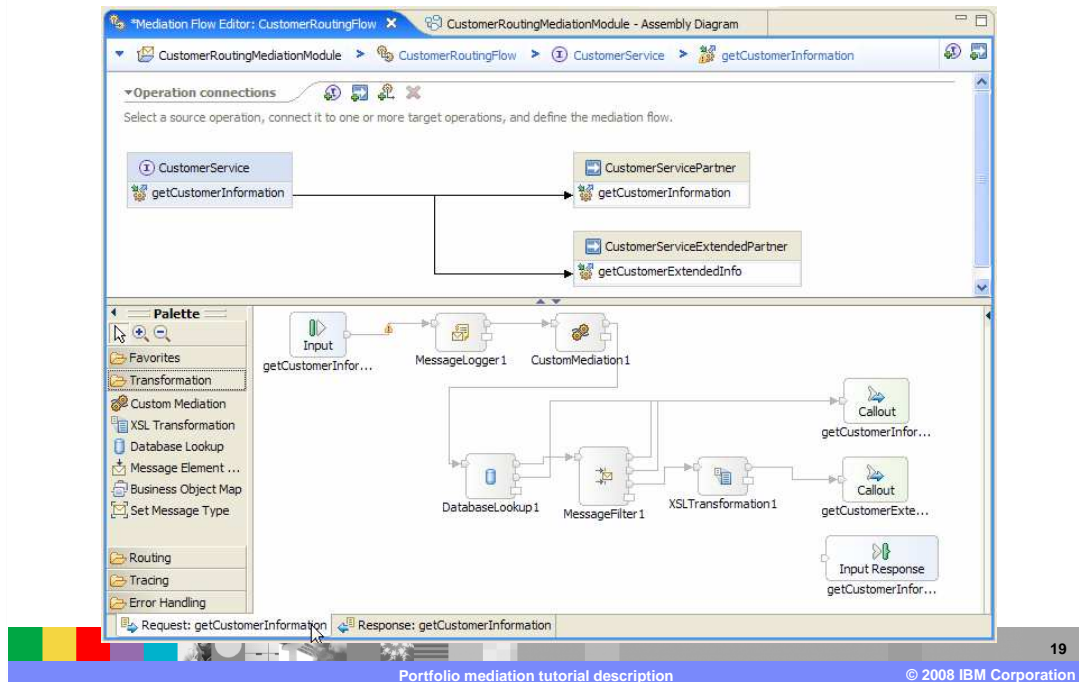
Portfolio mediation tutorial description

© 2008 IBM Corporation

Lab2 uses the results of Lab1 as a starting point.

First an Import is added to the assembly diagram, which supports the CustomerServiceExtended interface, as is shown in the screen capture. The mediation flow is enhanced to include the Custom Mediation, Database Lookup, Message Filter and XSLT primitives. There is an XSLT primitive also added for the response flow. These flows will be seen in screen captures on the next slide.

Tutorials – Lab2



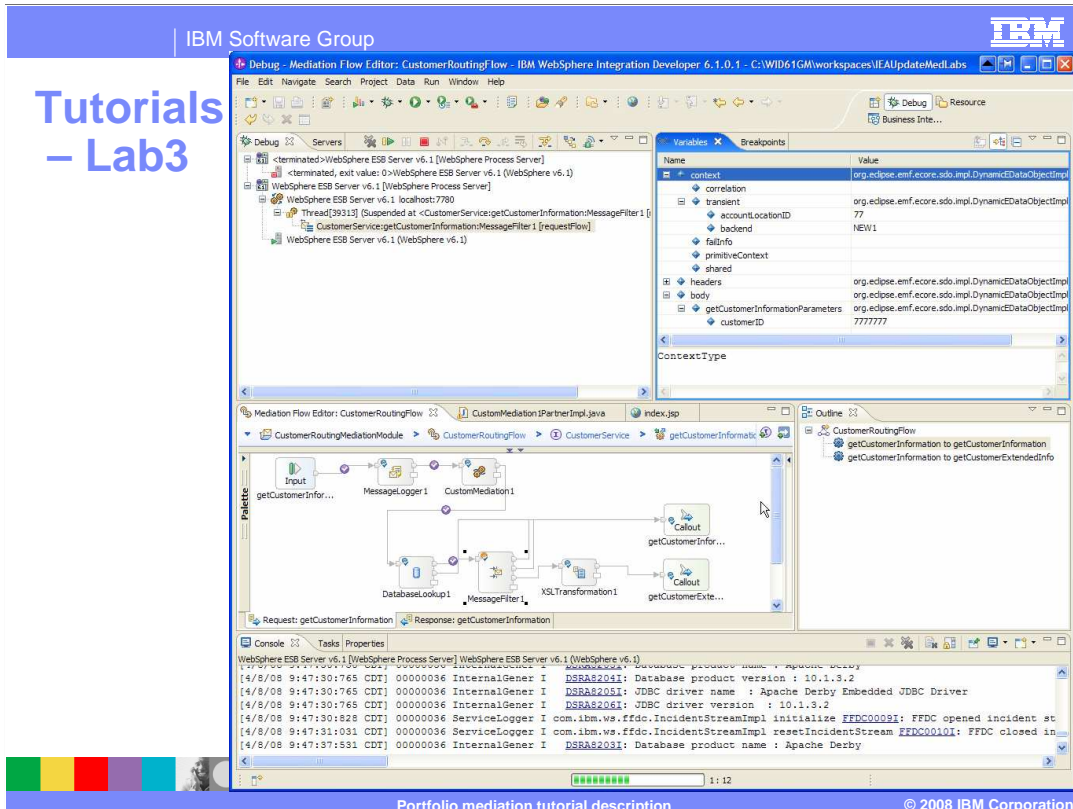
The mediation flow editor for Lab2 is shown here. In the Operation Connections panel you can see that the input operation on the CustomerService interface is wired to two possible output operations, one on the CustomerService interface and the other on the CustomerServiceExtended interface. The lower portion shows the canvas with the mediation flow which is identical to the flow described earlier.

Tutorials – Lab3

- Start with the results of lab2
- Use the visual debugger to examine the flow
 - ▶ Add breakpoints throughout the mediation flow
 - ▶ Deploy and run on server in debug mode
 - ▶ Examine the service message object (SMO) at various places in the flow



Lab3 starts using the end result from Lab2. In this tutorial you will use the Visual Debugger to examine the flow of a message through the mediation. This will be done by adding breakpoints on the mediation primitives, running the server in debug mode and examining the Service Message Object at various places in the flow. This will allow you to see the changes made to the Service Message Object by the Custom Mediation, Database Lookup and XSLT primitives.



In this slide you see a screen capture of the visual debugger. The panel in the middle shows the mediation flow, including the icons for breakpoints and icons on the wires showing through which wires the flow has already passed. The upper right panel shows the Service Message Object. You can see that the operation contains a customerID of all 7s. The transient context has been updated with the first two digits of this customerID and the backend has been identified as the new backend. The upper left panel shows exactly where you are in the flow and the lower panel is the console log from the server.

Tutorials – Test data

- The data used in the exercise is just hard coded in Java in:
 - ▶ CustomerBackend module
 - ▶ StockQuoteManager module
- The table shows the pre-defined customer IDs and portfolio values
- All other customer IDs return \$50

Customer ID	Backend	Portfolio Value
1111111	CustomerService (old)	\$16,850
2222222	CustomerService (old)	\$28,500
3333333	CustomerService (old)	\$6,880
4444444	CustomerService (old)	\$81,000
5555555	CustomerServiceExtended (new)	\$32,550
6666666	CustomerServiceExtended (new)	\$86,250
7777777	CustomerServiceExtended (new)	\$3,410
8888888	CustomerServiceExtended (new)	\$101,400
9999999	CustomerServiceExtended (new)	\$3,650

22

Portfolio mediation tutorial description

© 2008 IBM Corporation

Understanding the test data that has been built into the tutorials will help you predict and understand the behavior you will see when running the application using different customer IDs. The test data has been hard coded into the CustomerBackend module and the StockQuoteManager module. The only customer IDs that will be recognized are those shown in the table. The table also shows for each ID which backend will be used to process the request and the total portfolio value that will be computed. A customer ID that does not exist in this table will return a portfolio value of 50 dollars.

Tutorials – Test data

- The Database Lookup determines the backend to use
 - ▶ Prefixes directed to old backend: 11, 22, 33, 44
 - ▶ Prefixes directed to new backend: 55, 66, 77, 88, 99
 - ▶ A prefix not found in the database goes to the old backend through the keyNotFound terminal (for example: 14, 73, 42)
- How to determine which backend was used:
 - ▶ There is no way to tell from the browser interface
 - ▶ There are messages written to the console that show this:

```

O Entering getPortfolioValue() with 881111          O Entering getPortfolioValue() with 112345
O Getting customer info...                          O Getting customer info...
O Entering getCustomerExtendedInfo() with id: 881111 O Entering getCustomerInformation() with id: 112345
O Exiting getCustomerExtendedInfo() with id: 881111 O Exiting getCustomerInformation() with id: 112345
O Invoking getStockPrice() with SBC                O Invoking getStockPrice() with DELL
O Portfolio value is 50.0                           O Portfolio value is 50.0
O Value: 50.0                                       O Value: 50.0

```



23

Portfolio mediation tutorial description

© 2008 IBM Corporation

The database lookup determines which backend will be used. The database has been initialized so that customer IDs with prefixes of 11, 22, 33 or 44 will all go to the old backend. However, customer IDs with a prefix of 55, 66, 77, 88 or 99 all go to the new backend. Any other prefixes will not be recognized resulting in the keyNotFound terminal being used with the result of the flow going to the old backend.

Unfortunately, it is not possible to tell from the browser interface which of the backends processed the request. To determine that you must look at the output in the Console, specifically for the lines circled in red. On the left is output from the new backend showing the entry and exit of the getCustomerExtendedInfo operation whereas on the right is output from the old backend showing the entry and exit of the getCustomerInformation operation.

Summary

- The portfolio mediation scenario was described, including:
 - ▶ The original application before the mediation
 - ▶ Reasons for needing a mediation
 - ▶ The application containing the mediation
 - ▶ The flow logic within the mediation
- The tutorials were introduced
 - ▶ The scope of each of the tutorials was described
 - ▶ The built-in test data was explained



In this presentation the portfolio mediation scenario used by the tutorials was described. This included looking at the original application and then examining the business reasons for wanting to introduce a mediation. The application with the mediation and the mediation flow were then examined.

With that as a base, the individual tutorials were introduced and the test data used by them was also explained.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WPIv61_PortfolioMediationLabScenario.ppt

This module is also available in PDF format at:

..\WPIv61_PortfolioMediationLabScenario.pdf



You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

WebSphere

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.