



IBM Software Group

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

Augmentation, aggregation and retry tutorials



@business on demand.

© 2008 IBM Corporation
Updated May 21, 2008

This presentation introduces a series of tutorials designed to illustrate the mediation flow programming model for message augmentation, splitting and aggregation of messages and service invoke retry.

Goals and agenda

- Goal
 - ▶ Introduce the augmentation, aggregation and retry tutorials
 - ▶ Define capability each tutorial is designed to illustrate
 - ▶ Provide description of mediation flow logic and primitives used
 - ▶ Illustrate the flow using example data
- Agenda
 - ▶ Describe the scenario used for all flows
 - Input and output business objects illustrated with example data
 - Skeleton flow structure
 - ▶ Tutorial one – Augmentation
 - ▶ Tutorial two – Splitting and aggregating
 - ▶ Tutorial three – Fault recovery and retry
 - ▶ Tutorial four – Retry with alternate endpoints

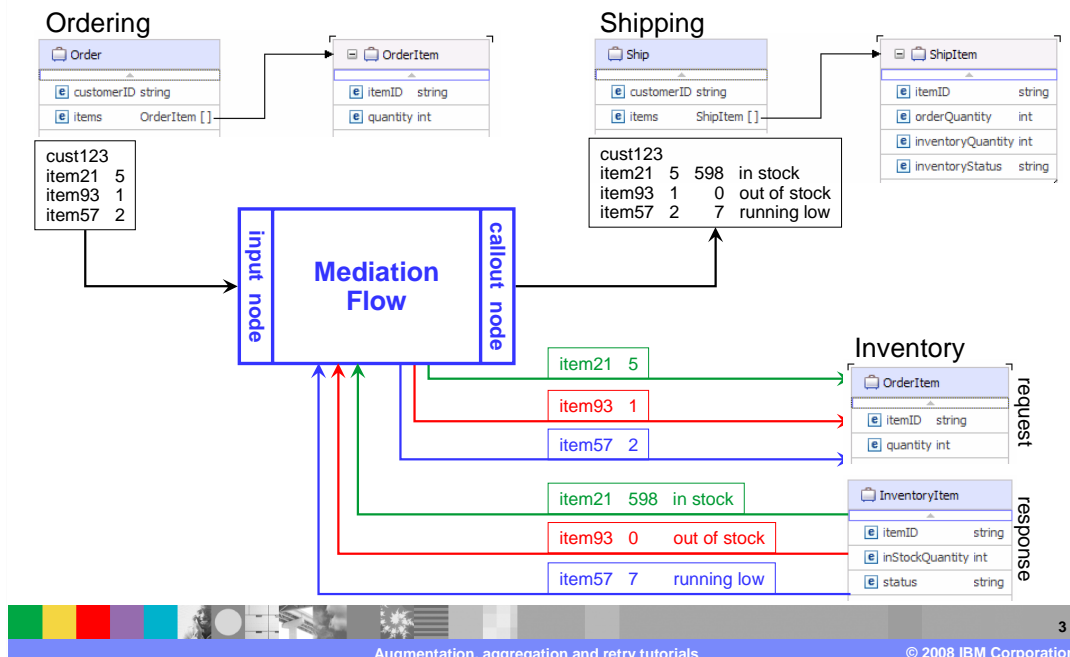


The goal of this presentation is to introduce to you a series of tutorials designed to illustrate elements of the mediation flow programming model. These include the use of the service invoke primitive to perform augmentation of data in a message. Also, the use of the fan out and fan in primitives to do message splitting and aggregation, enabling the processing of an array of elements within a message. Finally, the use of service invoke retry which performs automatic retry logic when a service call returns a fault. In the presentation there is a description of the capability each of the tutorials is designed to illustrate and how that is realized in the mediation flow logic. Example data is used to illustrate what the flow actually does.

The presentation starts out by describing the overall scenario that is used in all the tutorials. This is done by looking at the business objects used on input and output along with example data. The basic flow, or skeleton, used in the construction of each tutorial is described. Therefore, it does not have to be repeated in the description of each tutorial.

There are four tutorials, the first addressing message augmentation for a message with a single element. The next tutorial looks at how splitting and aggregating is done for augmenting a message with an array of repeating elements. The third tutorial looks at recovery from a fault, both in flow logic and with the use of service call retry. Finally, the fourth tutorial addresses the use of alternate endpoints for performing service call retry.

Overall flow scenario

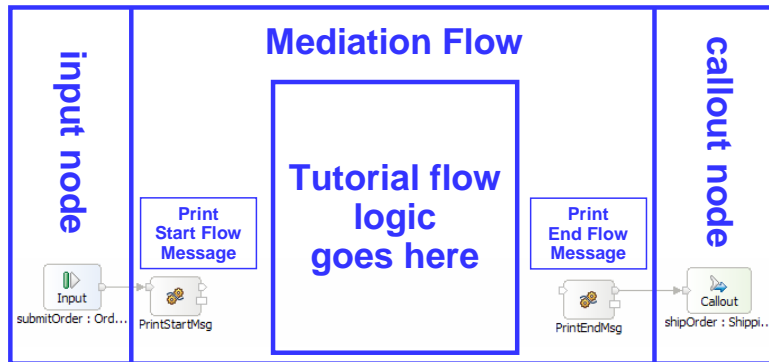


This slide provides you with a description of the overall scenario used in this series of tutorials. There is a mediation flow which takes as input a one way operation from an ordering system as shown in the upper left. It receives an Order business object instance containing a customerID field and an array of OrderItem business objects, each containing an itemID and quantity. You can see both the business object definition and example data on the slide. During the mediation, calls are made to an inventory service, one call for each item. The request passes an OrderItem as input and the response is an InventoryItem containing the itemID, the inStockQuantity and inventory status. This can be seen on the bottom of the slide. Finally, the mediation calls out to a shipping service in a one way operation, passing a Ship business object, as seen in the upper right of the slide. The Ship object is composed of a customerID and an array of ShipItem business objects, each containing the itemID, the orderQuantity, the inventoryQuantity and the inventoryStatus.

Each tutorial uses a variation of this basic scenario.

Skeleton structure of the flow

- tutorials use the same input node and callout node
- The flows start and end with primitives that write to the console (SystemOut.log)
 - ▶ The flow logic for all tutorials goes between these primitives
 - ▶ Screen captures in this presentation only show the tutorial flow logic




4

Augmentation, aggregation and retry tutorials

© 2008 IBM Corporation

Before looking at the individual tutorials, this slide introduces to you the basic skeleton flow within which the flow logic for each tutorial exists. The ordering and shipping services are the same throughout the tutorials, so that the endpoints of the flow are always the same and behave the same. Therefore, there are no changes to the input node or callout node. The tutorials make use of writing to the console view containing the SystemOut.log file as a way of illustrating what is happening in the overall processing. The skeleton flow contains a primitive immediately following the input node which writes to the console indicating the start of the flow. Likewise, there is a primitive just before the callout node which writes the end of flow message. The flow logic for all of the tutorials is contained between the two primitives used to write to the console. It is only this intervening logic that is shown in the illustrations of the tutorials.

Tutorial one – Augmentation



- Key illustration
 - ▶ Use of service invoke primitive for message augmentation
 - ▶ Requires use of XSL transformation and transient context
- Primitive usage
 - ▶ Order2Inventory XSL transformation
 - Modify body from order request message to inventory request message
 - Save order data in transient context
 - ▶ CheckInventory service invoke
 - Call inventory service to check item status
 - The results returned as an inventory response message
 - ▶ Inventory2Ship XSL transformation
 - Modify body from inventory response message to ship request message
 - Inventory data moved from body to body
 - Order data moved from transient context to body



The key purpose of tutorial one is to illustrate the use of the service invoke primitive. Calling out to a service from within a mediation flow can have many useful purposes. In this example it is used to obtain additional data that is then used to augment the contents of the message as it flows through the mediation. To perform a message augmentation scenario using a service invoke primitive, there is a need to transform the message to and from the message type used by the service invoke. This requires a primitive, such as the XSL transformation, that can modify the message type. It also typically requires use of the transient context to save information from the input message body that is needed by the callout message.

The logic for this flow requires three primitives as can be seen in the screen capture. The first is an XSL transformation, called Order2Inventory, whose primary purpose is to transform the message body from the order request message received by the mediation to an inventory request message needed to call the inventory service. It also places data into the transient context that needs to be saved across the call to the inventory service.


The second primitive is a service invoke, called CheckInventory, which makes a call to the inventory service which is external to the mediation flow. It passes information about the item being ordered and returns inventory information about that item.

The last primitive is another XSL transformation, called Inventory2Ship. Its purpose is to modify the message body from the inventory response message, returned by CheckInventory, into a ship request message, which is required by the callout. The body of the ship request message is created with data from two sources. Inventory information is taken from the body of the inventory response message and order information is taken from the saved data in the transient context.

Tutorial one – Augmentation

- Key assumption

- ▶ Input array contains exactly one item
- ▶ Array references are all indexed → [1]



6

Augmentation, aggregation and retry tutorials

© 2008 IBM Corporation

There is a key assumption made in the construction of this flow. Even though the input from order and the output to shipping are defined to have an array of items, this flow assumes that there is exactly one item in the input array containing the order. This requires, throughout the flow, using an index of one when dealing with the item arrays.

The illustration in the lower part of the slide shows the data flow associated with this mediation. The data coming from the order system is seen on the left. On the bottom of the illustration you see the data for the request and response with the inventory service. Finally, on the right, is the message to the shipping service. You can see that only one item is being processed and you can also see the augmentation of the data sent to the shipping service.

Tutorial two – Splitting and aggregating



■ Key illustration

- ▶ Use of fan out and fan in for splitting and aggregating
 - Enabling augmentation of multiple array elements
- ▶ Requires:
 - Augmentation scenario elements - service invoke, XSL transformation and transient context
 - Use of fan out context containing array element for each iteration
 - Use of shared context for aggregation of results
 - Use of message element setter append option for building up array

■ Primitive usage


- ▶ StartIteration fan out
 - Configured to iterate over item array in input message
 - SMO cloned at the start of each iteration
 - Fan out context set with current element at the start of each iteration



The key purpose of tutorial two is to illustrate the use of the fan out and fan in primitives to perform a splitting and aggregating scenario. Basically, this takes the previous augmentation scenario and enables the augmentation of multiple items in the message. Looking at the screen capture, you can see that the Order2Inventory, CheckInventory and Inventory2Ship primitives used in the previous tutorial are still part of the flow. To accomplish this scenario there are some additional requirements. The fan out context, which is initialized by the fan out primitive, is used to obtain the current item during each iteration. The shared context, which is a shared memory area, must be used to accumulate the results of each iteration. The message element setter append option is needed to build up the array in the shared context.

The best way to understand this flow is to walk through it to see what each primitive does in the flow. The first primitive is the StartIteration fan out primitive which is configured to iterate over the item array in the order request message. This primitive serves as the top of a loop in the flow, which is used to process the individual items. As each iteration is begun, the fan out primitive clones the SMO and sets the value of the current item into the fan out context. As a result of this, the SMO used to start each iteration is almost the same as the SMO in every other iteration. The only differences are that each SMO is in a separate memory area and the value of the fan out context is different.

Tutorial two – Splitting and aggregating



■ Primitive usage (continue)

- ▶ Order2Inventory XSL transformation (different map from tutorial one)
 - Modify body from order request message to inventory request message
 - Save order data in transient context
 - Customer data moved from the body
 - Item information moved from the fan out context
- ▶ CheckInventory service invoke (same as tutorial one)
 - Call inventory service to check item status
 - The results returned as an inventory response message
- ▶ SaveInventoryInfo message element setter
 - Inventory data moved from body to transient context (combined with order data)



The next primitive is the Order2Inventory XSL transformation, which has the same purpose as its counterpart in the previous tutorial, but contains a different mapping. It modifies the body of the message from the order request message to the inventory request message, in preparation for calling the inventory service. Information about the order is saved in the transient context, with the customer data coming from the body of the order message and the item information coming from the fan out context. The order information saved in the transient context is placed into a Shipltem structure in preparation for being combined with inventory information.

The next primitive, CheckInventory, is unchanged from the previous tutorial and is used to call the inventory service.

The next primitive in the flow is a message element setter named SaveInventoryInfo. It takes inventory information about the item and moves it from the body of the inventory response message to the transient context. It is placed into the Shipltem structure in the transient context, thus combining it with the order data for the item already placed there.

Tutorial two – Splitting and aggregating



■ Primitive usage (continue)

- ▶ AggregateItemInfo message element setter
 - Order and inventory combined data for this item taken from transient context and appended to array in shared context
- ▶ EndIteration fan in
 - Configured to complete when input array fully processed
 - Flow returns to StartIteration fan out if not complete
- ▶ Inventory2Ship XSL transformation (different map from tutorial one)
 - Modify body from inventory response message to ship request message
 - Customer information moved from transient context to body
 - Completed item array with order and inventory information moved from shared context to body



Continuing in the flow, another message element setter appears next, which is named `AggregateItemInfo`. This is used to move the completed `ShiplItem` from the transient context and append it to the array of `ShiplItem` elements in the shared context. The shared context is in a memory area shared by all of the SMO clones, and therefore this provides a mechanism for accumulating the results of each iteration.

The next primitive is the fan in, called `EndIteration`. This primitive serves as the end of the iterative loop. It has configuration information defining its completion criteria, indicating that it completes when an entire item array has been processed. If the item array processing is not complete, the flow returns to the `StartIteration` fan out for the next iteration. If the item array processing is complete, the flow continues from the fan in.

When the fan in completes, the flow proceeds to the `Inventory2Ship` XSL transformation. Similar to the previous tutorial, this primitive modifies the message so that it can be sent to the shipping service. However, the actual mapping required is different. The customer information is moved from the transient context, where it was saved at the beginning of the flow, to the body. The array of ship items, which as been built up in the shared context during the iterative flow, is moved from the shared context to the body.

Tutorial two – Splitting and aggregating



- Best practices for handling of array elements in aggregation
 - ▶ Target array element type is used throughout the flow
 - Use this as a guide to design the transient and shared contexts you use
 - ▶ Array type in shared context is identical to the target array type
 - ▶ Have single element in the transient context
 - ▶ Build up single element during an iteration
 - Use whatever primitives are needed for your scenario
 - For example, in this scenario an XSL transformation and message element setter
 - ▶ Append single element to shared context array at end of an iteration
 - ▶ Move shared context array to message body array when all iterations complete



10

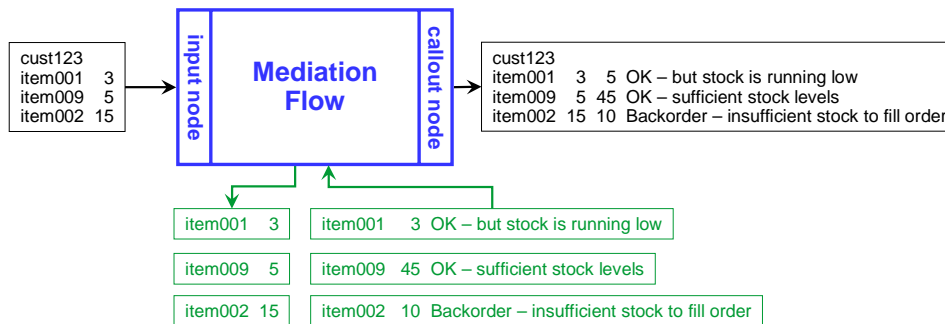
Augmentation, aggregation and retry tutorials

© 2008 IBM Corporation

This slide is a tangent to the specific tutorial discussion, but is inserted here to highlight key factors about the processing done in this tutorial. These points are worth mentioning as they are considered best practices for implementation of an aggregation scenario.


The scenario assumes there is a target array of repeating elements that is being constructed during the aggregation. You should use the array element type of the target array when defining your shared and transient contexts for use in the flow. The array defined in the shared context should be the same type as the target array, composed of the same element types. In the transient context, define a single element of the type contained in the array. During the iterative flow, build up the contents of the single element in the transient context, using whatever primitives are needed for your particular scenario. For example, in this scenario, an XSL transformation was used to move in the order information and a message element setter was used to move in the inventory information. At the end of each iteration, use a message element setter to append the single element built in the transient context to the end of the array being built up in the shared context. When the entire input array has been processed, move the array built up in the shared context to the message body.

Tutorial two – Splitting and aggregating



This slide shows an example of the processing that occurs as a result of this flow, assuming the input order contains three items. This is shown on the left. The three inventory service requests and responses, one for each item, can be seen on the bottom. The callout to the shipping service, with augmented data for all three items, appears on the right.

Tutorial three (first section) – Fault recovery



▪ Key illustration

▶ Recover from a fault returned to the flow

- Allows full array to be processed regardless of the fault returned from the inventory service for one item

▪ Primitive usage

▶ Same flow as tutorial two with addition of one primitive

▶ FaultRecovery XSL transformation

- Wired between the CheckInventory service invoke primitive's terminal for InventoryFault and the AggregateItemInfo message element setter
- Converts message from inventory fault message to inventory response message
- Sets inventory quantity to 0 and inventory status to indicate an error occurred
- Other than actual inventory data values the entry to AggregateItemInfo appears as if the call worked




The third tutorial is broken into two sections, the first for fault recovery and the second illustrates service call retry.

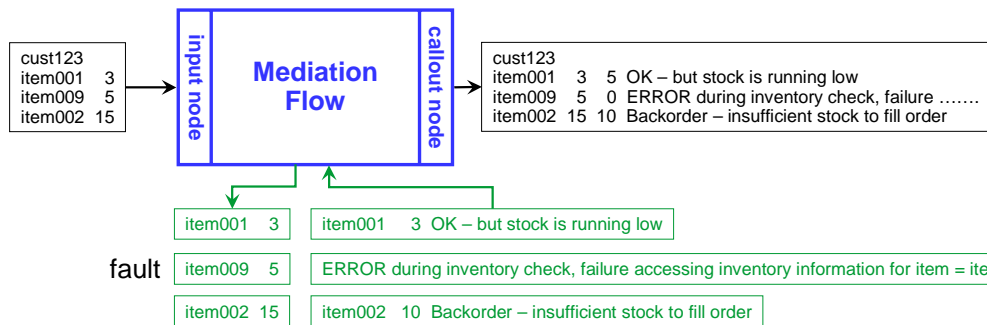
This first section describes fault recovery in the flow, allowing the flow to continue rather than terminate when a fault is returned from the inventory service. This enables the entire list of items to be passed to the shipping service, but without inventory information for the item being processed when the fault occurred.

To do this the same flow is used as tutorial two, with the addition of one primitive, the FaultRecovery XSL transformation. The InventoryFault terminal, of the CheckInventory service invoke primitive, is wired to this XSL transformation. It converts the SMO from an inventory fault message to an inventory response message. The inventory quantity is explicitly set to zero and the inventory status is set to indicate that an error occurred calling the inventory service. By performing this transformation, the SMO appears as if the inventory service call worked, with the exception of the actual data values for inventory quantity and status. The FaultRecovery XSL transformation is then wired to the AggregateItemInfo and processing proceeds as if no fault occurred.

Tutorial three (first section) – Fault recovery




- Inventory service
 - ▶ Implementation randomly works or returns a fault
 - ▶ Because random, your output will vary



When doing this part of the tutorial, the endpoint for the inventory service is changed to an implementation that is designed so that it sometimes works and sometimes returns a fault. Because the results are random, the output you see when testing can be different than what is illustrated in this example. In the example, the input is the same three items used in the previous tutorial. Notice that the call to the inventory service for the second item, item009, returned a fault. The output sent to the shipping service shows all the items in the list but with the inventory information for item009 indicating a problem calling the inventory service.

Tutorial three (second section) – Retry



- Key illustration

- ▶ Service invoke automatically retries when fault returned
 - Reduces chance of fault actually being returned to the flow

- Primitive usage

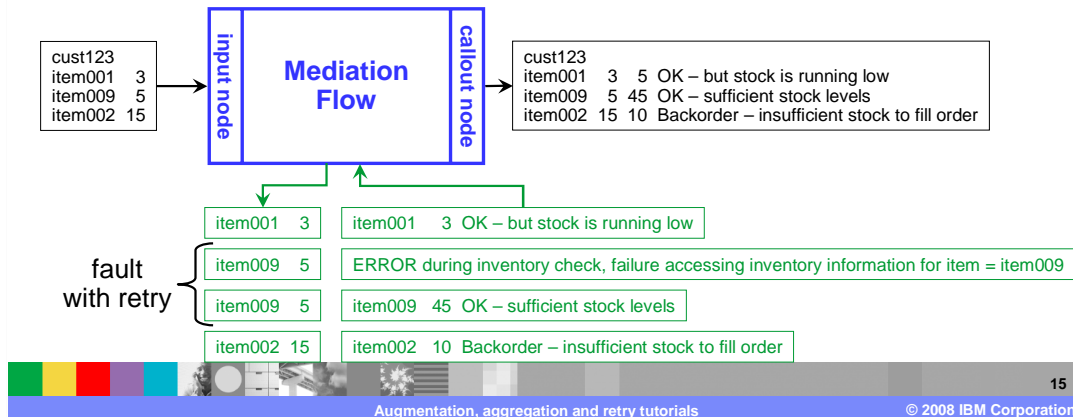
- ▶ No changes to flow
- ▶ CheckInventory service invoke primitive
 - Configured for retry on modeled fault with retry up to three times



The second section of tutorial three illustrates the use of the service call retry capabilities of the service invoke primitive. By performing a retry automatically, the likelihood of the service invoke primitive having to return a fault is reduced because the retry might be successful. To do this, there are no changes made to the flow other than the configuration for the CheckInventory service invoke primitive. It is modified so that it performs a retry on the occurrence of a modeled fault, performing up to three retries. If any retry is successful, the flow continues. If not, the fault returned on the third retry is passed to the flow.

Tutorial three (second section) – Retry

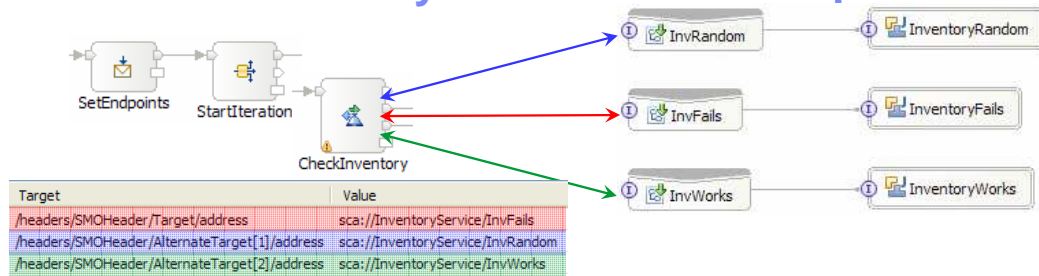
- Expected behavior in the tutorial
 - ▶ Inventory service randomly works or returns a fault
 - ▶ Inventory service will not return more than two sequential faults
 - ▶ Inventory service will not work more than two time sequentially
 - ▶ Retry count on service invoke is set to three (result is the fault terminal is never used in this tutorial)



In this tutorial, the inventory service that randomly works is being used. It is implemented so that it will never return more than two sequential faults, nor will it work more than two times in a row. Therefore, the results are not entirely random. Consequently, when you are running the test, at least one fault will occur. Also, the retry limit is never reached, so the inventory calls for every item will always eventually be successful.

In the example data for the flow, the same three order items previously used are being passed into the mediation. Notice that there are actually four calls made to the inventory service rather than three. The call for the second item, item009, returned a fault. This resulted in the service invoke primitive performing a retry which was successful. In the output data sent to the shipping service you can see that inventory information has been included for all three items. The only evidence that a fault occurred is the extra call to the inventory service, which was only seen by the CheckInventory service invoke primitive.

Tutorial four – Retry with alternate endpoints



- Key illustration
 - Service invoke uses alternate endpoints when performing retry
- Primitive usage
 - SetEndpoints message element setter
 - Added at beginning of the flow to initialize the target and alternate target URLs
 - CheckInventory service invoke primitive
 - Configured to use dynamic target address on initial call
 - Configured for retry on modeled fault with retry up to three times
 - Configured to use alternate addresses on retry

16

Augmentation, aggregation and retry tutorials

© 2008 IBM Corporation

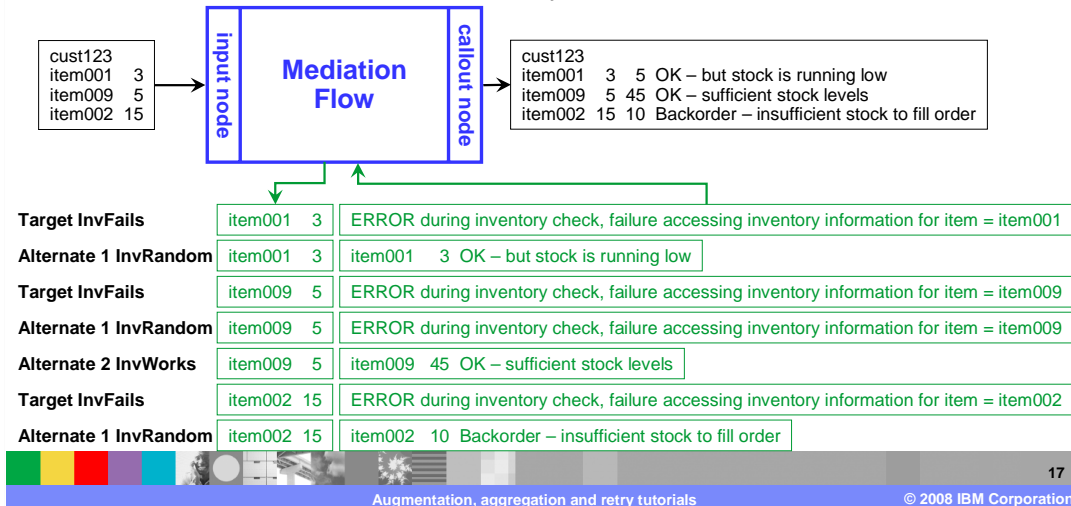
Tutorial four illustrates the use of alternate endpoints when doing a service call retry. To do this, the target address field and alternate target address array in the SMO must be initialized with appropriate endpoint URLs. The initialization of these fields is typically accomplished using the endpoint lookup primitive which calls the WebSphere Service Registry and Repository. However, in this tutorial the fields are set using a message element setter.

The SetEndpoints message element setter is added to the flow before the StartIteration fan out primitive. It sets the SMO target address field and two alternate target addresses using URLs that point to exports in the module containing the inventory service. The CheckInventory service invoke is modified so that it uses the target address URL for the initial call, and then uses the alternate target addresses for the retries. In the screen captures, you can see the URLs for the targets set by the message element setter. You can also see the CheckInventory service invoke and the assembly diagram for the inventory service containing the three exports which are the service endpoints.

Tutorial four – Retry with alternate endpoints

Understanding the endpoints

- ▶ Target InvFails always returns fault
- ▶ Alternate 1 InvRandom sometime returns fault, sometimes works
- ▶ Alternate 2 InvWorks always works



The three endpoints used in the tutorial are implemented so that one will always return a fault, one will randomly work or return a fault and the third will always work. To aid in illustration of this function, the target address is set to the endpoint that will always fail. The first alternate target is set to the endpoint that will randomly work and the second alternate target is set to the endpoint that always works. As a result, the first call to the inventory service for each item will always result in a fault. The first retry may be successful or return a fault. If the first retry returns a fault, the second retry is performed and it will always work.

Looking at the data, the same three order items are passed in. You can see seven calls to the inventory service. For item001 and item002, there is one failing call followed by success on the first retry. For item009, the first call fails as does the first retry, but the second retry works. Looking at the output going to the shipping service, you can see that the item array is fully populated with inventory information. The only evidence that faults occurred are the extra calls to the inventory service, which were only seen by the CheckInventory service invoke primitive.

Summary

- Introduced the augmentation, aggregation and retry tutorials
- Defined the capability each tutorial is designed to illustrate
- Provided a description of the mediation flow logic and the primitives used
- Illustrated the flows using example data



In summary, this presentation examined a series of tutorials. The tutorials are designed to illustrate message augmentation, message aggregation and service call retry. Each of the four tutorials was described in terms of the capability the tutorial is designed to illustrate. The mediation flow logic and primitives used were described. Finally, for each tutorial, example data was used to illustrate the behavior of the scenario.

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_WBIV61_61ProgModellAbsIntro.ppt

This module is also available in PDF format at: ..\\WBIV61_61ProgModellAbsIntro.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:

IBM WebSphere

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

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 NON-INFRINGEMENT. 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.