



IBM Software Group

IBM® WebSphere® Application Server V6

Enterprise JavaBeans™ (EJB) Mediator



@business on demand.

© 2004 IBM Corporation
Updated January 25, 2005

This presentation, will focus on providing an overview of the EJB Mediator.

Goals

- Provide an overview of the EJB Mediator
- Describe two modes of operation



The primary goal of this presentation is to provide an overview of the EJB Mediator.

EJB Mediator: Introduction

- Provides SDO Data Mediator for applications implemented with Entity EJBs
- Intended for Java programmers who are familiar with
 - ▶ J2EE programming model
 - ▶ EJB Query Language
- Ideal choice when application is implemented with Entity EJBs and needs to meet any of the following requirements
 - ▶ Return an XML document from an EJB application
 - Service oriented architecture
 - ▶ Caller needs disconnected, serializable data
 - EJBs are shared, transactional objects
 - EJBs are not serializable (remote proxy)

The EJB Mediator provides an SDO data mediator for applications developed using Entity EJBs (CMP type only). The EJB Mediator allows programmers to specify a compound EJB query statement to access entity EJB data through the DataGraph/DataObject interfaces. The EJB Mediator is targeted primarily for Java programmers that are already familiar with the J2EE programming model and the EJB Query Language.

A natural question is, when would you use the EJB Mediator? The EJB mediator should be considered when an application is architected to use entity EJBs and any of the following items applies to the application design:

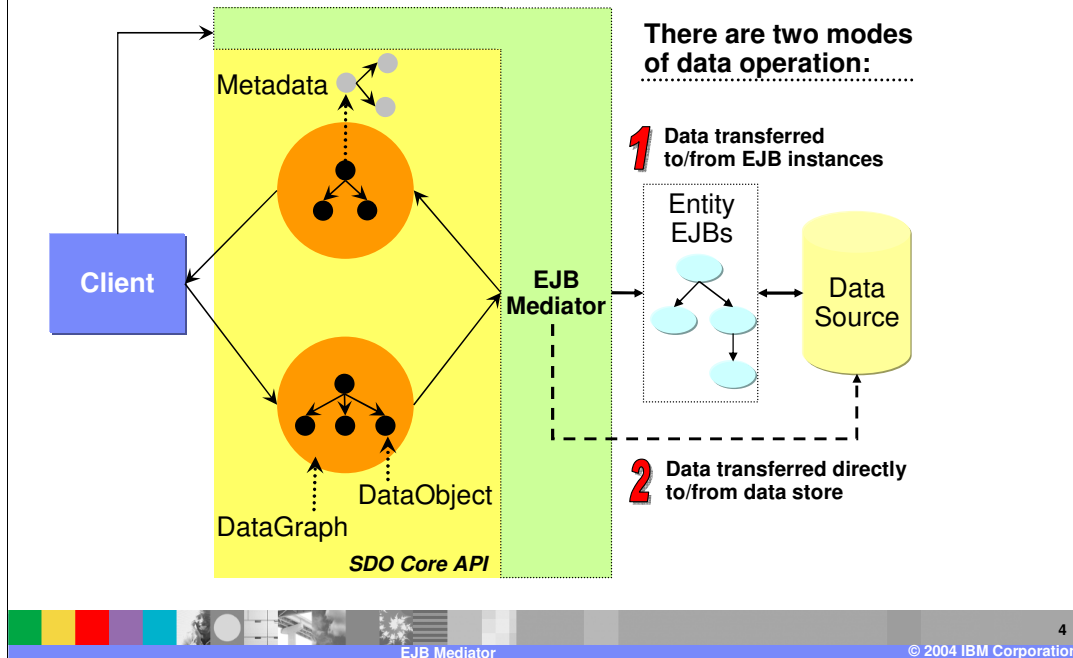
--> You need to return an XML document from an EJB application (for example, in the case of a service oriented architecture)

--> Caller needs disconnected, serializable data

--> Application is being built using tools that use DataObjects (for example, IBM Rational Application Developer V6 will have new EJB generation tools that will replace the EJB Access bean APIs with DataObjects and Mediators)

One final implementation note: Although the EJB Mediator facilitates access to entity EJBs, it is not itself an EJB.

EJB Mediator: Big Picture



From the client perspective, the EJB mediator looks very similar to any other type of data mediator. Recall from previous discussions that all mediators provide a uniform view of the backend data through the SDO core APIs (DataObject and DataGraph). However, there are typically a small set of features that are specific to a particular type of mediator. The diagram shown on this slide is intended to highlight one of the important features provided with the EJB Mediator.

When using the EJB mediator, there are two modes of operation to work under when accessing entity EJB data. In the first mode of operation, the EJB mediator uses active EJB instances to transfer data to/from the data source. However, in the second mode of operation, the EJB mediator transfers data to and from the data store directly rather than activating the EJB.


EJB Mediator: Data Graph Example

- EJB Mediator constructs a Data Graph based upon interrelated EJB query statements
 - ▶ EJB CMP fields and EJB method return values make up the SDO attributes
 - ▶ EJB CMR values make up SDO references

EXAMPLE

```
select d.id, d.name from
DeptBean d
  where d.name like 'p%'
select e.id, e.salary from in
(d.emps) e
  where e.salary > 50
select p.* from in(d.projs) p
select t.* from in(p.tasks) t
```

Data Graph



EJB Mediator

5

© 2004 IBM Corporation

The EJB mediator allows a client to specify an EJB query statement that is used to build a data graph. The attributes for the DataObjects making up the DataGraph correspond to the EJB CMP fields or return values from EJB methods included in the EJB query statement. The SDO references (the way in which a DataObject in the graph refers to another DataObject) are related to EJB CMR values.

Included on this slide is an example EJB Query statement, along with visual representation of the DataGraph that would result from this query. The following lists the schema for this example:

---- EJB Schema ----

Department

CMP Fields (id, name, budget)

CMR (emps=>Collection of Employee, projs=>Collection of Project, mgr=>Employee)

Employee

CMP Fields (id, name, salary, bonus)

CMR (dept=>Department, tasks=>Collection of Task, manages=>Collection of Department)

Project

CMP Fields (id, name, cost)

CMR (tasks=>Collection of Task, dept=>Department)

Task

CMP Fields (id, name, desc)

CMR (proj=>Project, emps=>Collection of Employees)

EJB Mediator: High-level Function

Functions	Description
Create EJB Mediator	Client provides a compound EJB query statement
Get a Data Graph	Mediator queries backend with EJB query statement and creates Data Graph from results
Apply changes	Mediator updates changes from Data Graph to the data source
Get schema	Mediator provides schema information associated with the Data Graph



The high level functions for the EJB Mediator are similar to the JDBC Mediator. These high level functions represent the functions that most mediators would be expected to support.

Note that the “Get Data Graph” and “Apply changes” functions can be applied directly to the database or through activated EJB instances as discussed previously.

EJB Mediator: Design Points

- Creating an EJB Mediator
 - ▶ Client specifies compound query statement by providing an array of strings
 - ▶ Query syntax maps to EJB Query Language as closely as possible
- Get Data Graph
 - ▶ Data Graphs can be constructed without activating EJBs
 - ▶ Queries can be specified to search only EJBs that are activated or in the cache



This slide and the next highlights several design points for the EJB Mediator.

When creating an instance of the EJB Mediator, the client needs to specify a compound query statement. This is done by providing an array of strings that represent the individual queries that make up the compound EJB query statement. It is important to note that every effort has been made to map the syntax of the compound query statement passed in to create the EJB mediator to the EJB Query Language. However, there are several modifications included to support various functions of the EJB Mediator. An example where the EJB QL syntax has been augmented is with respect to using the EJB mediator to only query EJBs from a user supplied collection of activated EJBs.

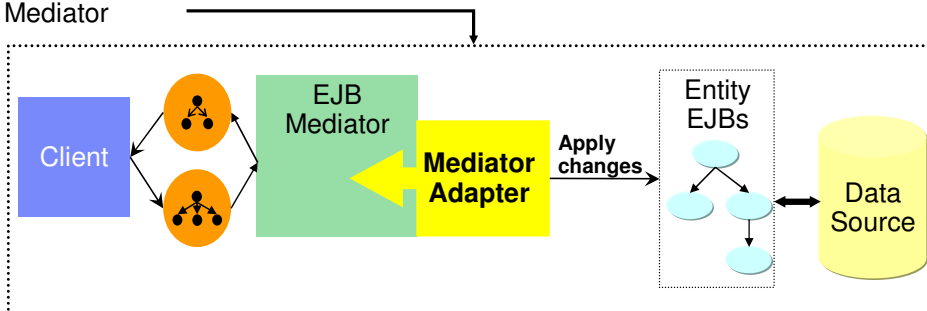
One of the design features of the EJB Mediator is that for the get Data Graph function the client has a choice to have the data graph constructed without activating the EJBs. In this case, the mediator accesses the data store directly to build the data graph rather than activating the EJBs. In addition to this, the client can specify to the EJB Mediator only to search EJBs that are currently activated or are in the cache. The advantage to these features is that avoiding EJB activation can improve performance.

EJB Mediator: Design Points (Continued)

Apply Changes

Several choices for updating data objects

- Changes translated to SQL and data source updated directly by EJB Mediator
- Changes written back through EJB by EJB Mediator
- Client may specify mediator adapter to handle updates on behalf of EJB Mediator



When changes to the data graph are written back to the data store, there are several choices. The first choice is to let the EJB mediator handle the updates completely. In this case, the changes are translated to SQL and the data source is updated directly. The next choice is to have the mediator write back the changes through the EJBs. And finally, the last choice allows the client to provide a mediator adapter that allows the client to handle the updates on behalf of the mediator. These updates are written back through the EJBs. This is a pluggable interface that is implemented by the client to handle the updates.

Note: The mediator adapter can handle creates, updates, and deletes on behalf of the mediator. For more information on the mediator adapter see the com.ibm.websphere.sdo.mediator.ejb.MediatorAdapter interface.

EJB Mediator: Entity EJB Requirements

- EJBs must have an EJB Local interface
- The following must be promoted to the local interface:
 - ▶ Get/Set methods for CMP fields
 - ▶ EJB methods used in query statements
- The following must be included on the EJB Home
 - ▶ create(PrimaryKeyClass)
 - Alternatively, you can supply a mediator adapter to handle the create
 - ▶ findByPrimaryKey(PrimaryKeyClass)
 - ▶ remove(Object key)



Because the EJB Mediator uses local interfaces for EJBs there are several important points to remember when using this mediator. First, getter and setter methods for CMP fields or EJB methods included in compound query statements used by the mediator must be promoted to the local interface.

Also, if the EJB Mediator is used by the client to create an EJB, the client must be sure that the create(Primary_Key) method is included on the EJB Home. Alternatively, if this is not the case, you may supply a mediator adapter to handle the EJB create on behalf of the EJB Mediator.

SDO: Summary

- Provides a SDO data mediator for EJB data sources
- Integrated with IBM Rational Application Developer V6
 - ▶ Create Session Bean Façade



In this presentation you learned about the SDO-based EJB Mediator. This technology is a new feature in the WebSphere Application Server V6 runtime environment and is also supported in IBM Rational Application Developer V6. Specifically, the EJB Mediator functionality is available as the Create Session Bean Façade from the context menu for Entity EJBs. The resulting session bean can be added to a JSF page from a Session Bean data source from the Page Data view.

References

- Developer Works:
 - ▶ <http://www.ibm.com/developerworks/library/j-commonj-sdowmt/>

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	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AW	DB2 Universal Database	Lotus	pSeries	zSeries

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

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

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

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 and/or changes in the product(s) and/or program(s) 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 (e.g., 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 2004. 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.