

Java/EJB Translation Rules

version 0.6
6/10/02

copyright 2002 Pathfinder Solutions, all right reserved.



1.	INTRODUCTION	3
2.	REFERENCED DOCUMENTS	4
3.	TERMS AND DEFINITIONS	4
4.	TARGET DOCUMENT SET	4
4.1	DIRECTORY HIERARCHY	5
4.2	JAVA FILES	5
4.3	RELATIONAL DATABASE SCHEMAS	5
4.4	EJB SPECIFIC FILES	5
4.5	MECHANISMS	5
5.	STRUCTURAL DESIGN	5
5.1	HARDWARE	6
5.2	SOFTWARE	6
6.	MBSE TO RELATIONAL DATABASE TRANSLATION RULES	6
6.1	ATTRIBUTE	6
6.2	OBJECT (UML CLASS)	6
6.3	RELATIONSHIP	6
6.3.1	<i>BinaryRel</i>	6
6.3.2	<i>SubSuperRel</i>	7
6.4	CURRENT STATE	7
6.5	INTERDOMAIN RELATIONSHIPS	7
7.	MBSE TO JAVA/EJB TRANSLATION RULES	8
7.1	DOMAIN CHART CONSTRUCTS	8
7.1.1	<i>Bridge</i>	8
7.1.2	<i>Constant</i>	8
7.1.3	<i>DataType</i>	8
7.1.4	<i>UserDefinedType</i>	9
7.1.5	<i>Domain</i>	9
7.1.6	<i>System</i>	10
7.2	INFORMATION MODELING CONSTRUCTS	10
7.2.1	<i>Attribute</i>	10
7.2.2	<i>Object (UML Class)</i>	11
7.2.3	<i>Relationship</i>	11
7.3	STATE AND SERVICE MODELING CONSTRUCTS	12
7.3.1	<i>Action</i>	12
7.3.2	<i>Event</i>	13
7.3.3	<i>Generate</i>	14
7.3.4	<i>State Transition Table</i>	14
7.3.5	<i>Output Parameters</i>	14
7.4	ACTION LANGUAGE	14
7.4.1	<i>AttributeSelection</i>	14
7.4.2	<i>Create</i>	15
7.4.3	<i>CreateServiceHandle</i>	15
7.4.4	<i>Delete</i>	15
7.4.5	<i>Expression</i>	15
7.4.6	<i>Instance Lookups - Find/Foreach</i>	15
7.4.7	<i>InvokeServiceHandle</i>	16
7.4.8	<i>Link</i>	16
7.4.9	<i>Navigation</i>	16
7.4.10	<i>ServiceInvocation</i>	16
7.4.11	<i>SubSuperNavigation</i>	16
7.4.12	<i>Unlink</i>	16
7.5	MBSE SEMANTICS	17
7.5.1	<i>Timers</i>	17
7.5.2	<i>Instance Lists</i>	17

7.5.3	<i>Service Handles</i>	17	
7.5.4	<i>Task</i>	17	
7.6	JAVA SPECIFIC ISSUES	18	
7.6.1	<i>Exceptions</i>	18	
7.7	EJB SPECIFIC ISSUES	18	
7.7.1	<i>Loopbacks</i>	18	
7.7.2	<i>Exceptions</i>	19	
7.7.3	<i>Transactions</i>	19	
7.8	REALIZED INTERFACES	20	
8.	MODELING CONVENTIONS AND RESTRICTIONS	20	

Version	Date	Person	Comments
0.1	6/15/00	Greg Eakman	Created
0.2	6/20/00	Greg Eakman	Feedback from Tom M and Steve R
0.3	6/29/00	Carolyn Duby	Fix several typographical errors
0.4	7/3/00	Greg Eakman	Update with Carolyn's comments, add transactions and exceptions
0.5	7/17/00	Greg Eakman	Add comments
0.6	7/10/02	Greg Eakman	Clean up

1. INTRODUCTION

Model Based Software Engineering (MBSE), as presented by Pathfinder Solutions, is a rigorous, complete, unambiguous, high level executable model of a particular problem domain. As such, this model can be translated into other executable implementations in the form of source code.

Model data is stored in an MBSE meta-model database. A set of translation rules in the form of ASCII templates, read and interpreted by Springboard, produces the source code. Springboard is a translation engine that extracts the semantic information contained in your UML analysis models and presents it in textual form via a flexible and simple template notation. See the Springboard Users Guide for more information on what data is accessible through the templates and the format and structure of templates.

This document describes the mapping of analysis into a Java/EJB design using Springboard and translation templates. The reader is assumed to have a working knowledge of Model Based Software Engineering (MBSE). The reader is also assumed to be familiar with the Unified Modeling Language (UML™) and EJB. See Enterprise JavaBeans, Second Edition, Richard Monson-Haefel, published by O'Reilly and associates, for more information on EJB.

It is also assumed the reader is familiar with the terminology, concepts and conventions presented in:

"Model Based Software Engineering – an Overview of Rigorous and Effective Software Development Using UML™" 1998, Pathfinder Solutions, Inc.

The document is organized by the major constructs within the MBSE meta-model. A brief description of each analysis construct is provided, followed by the mapping rules used to generate Java source code to be integrated with an EJB container and server. Where choices in the implementation can be made, properties are defined which guide the templates on how to translate the construct. These properties can be set either within the Model Editor through Pathfinder Extensions, or can be set at translation time using the SetProperty capability in the Springboard Template Language. Properties can also be inferred by the translation templates based on the models themselves.

2. Referenced Documents

Enterprise JavaBeans, Second Edition, Richard Monson-Haefel, O'Reilly and Associates

Springboard User's Guide

"Model Based Software Engineering – an Overview of Rigorous and Effective Software Development Using UML™" 1998, Pathfinder Solutions, Inc.

3. TERMS and DEFINITIONS

Templates a set of text files that capture patterns to define the structure, fixed contents, and variable contents (corresponding to OOA model elements) of the target document set, using Syntactic Elements that conform to the notation specified in the Springboard User's Guide.

EJB Enterprise Java Beans

Extraction the portion of the translation process when Springboard reads the OOA modeling information captured in your CASE database

JVM Java Virtual Machine. This is the runtime system responsible for the execution of Java bytecodes produced by compiling Java source code.

MBSE Meta-Model the MBSE/UML Information Model that describes MBSE and stores the application models for translation.

Parsing the portion of the translation process when Springboard reads the templates in preparation for production

Production the portion of the translation process when Springboard executes the templates, populating your target document set with the OOA information acquired during extraction

Property a name-value pair associated with a MBSE/UML model element

Target document a file or document produced by Springboard as the result of executing an OUTFILE directive

Target document set the complete set of target documents produced by a single execution of Springboard on a set of templates

Translation the execution of Springboard performs 3 steps in order: Template Parsing, OOA model Extraction, and Target Document Production

4. Target Document Set

Considering the task of translation at the highest level, the goal is to establish a set of instances of target documents with a set of OOA analysis data according to specified templates, or templates. The templates specify the specific analysis information required, how it's arranged, and in what context.

The above paragraph helps us identify the major tasks/components of a translation:

- Develop a set of templates to define our target document set
- Employ Springboard to validate the correctness of the new templates
- Apply the proven templates with Springboard to produce the Target Document Set

The target document set consists of a set of Java files, organized in a directory structure that reflects the hierarchical Java package structures. The target document set also includes, where possible, deployment descriptors for the target EJB server. The target document set also includes the relational database schemas derived from the models.

4.1 Directory Hierarchy

Java organizes classes by packages, and, in the development environment, these package names are reflected in the directory structure. The generated Java files will be placed in a directory structure based on the Domain Chart, with the system name (Domain Chart name) as the top level class, and each of the domains as a subdirectory and subpackage of the system.

A Package property on the domain chart will be available to specify the upper parts of the package hierarchy. Most packages reflect the DNS domain name of the company producing the Java software – e.g., com.pathfindersol.java.mechanisms. The corresponding high level directories will be created in the document set.

4.2 Java Files

Generally, each UML class and Domain will result in one Java file that contains the implementation of the class. In some cases, more than one file and Java class is required to support the EJB design.

Java filenames must reflect the top level class they contain.

The content of the files, the Java source code, will conform to de-facto standard format and naming conventions, and will be tailored to any coding standards required by the client.

In addition, a makefile is generated for the package that uses the Sun JDK command line compiler. If a standard IDE is chosen as the Java development environment, it may be possible to generate the project file for the IDE.

4.3 Relational Database Schemas

The relational database schemas consist of a set of instructions to create tables and columns within a relational database. The schemas use basic SQL constructs to construct the database.

During development, it is expected that the Information Models, and the derived database schema will change often. Depending on the scope of the changes, the existing tables can be destroyed and wholly replaced by the new set of tables. A process must be designed by which existing data in the database is not lost in this process. One alternative is to back up the data before clearing the database, then restoring it into the new tables. Another approach is to pass the new tables through a diff-like filter that can compare the old and new schemas, then generate appropriate SQL commands to update the existing table, preserving all data.

4.4 EJB Specific Files

There are some EJB specific files that can be generated based on the UML models and properties attached to them. The XML deployment descriptor for EJB 1.1 is one of them. Depending on the EJB server specific deployment process, other files may be generated to aid and automate the deployment process.

4.5 Mechanisms

A set of mechanisms, in the form of source code, will be provided by Pathfinder Solutions. These mechanisms provide the classes required to execute the MBSE semantics within the Java language.

5. Structural Design

This document mainly addresses the Mechanical Design aspect of system generation, translation rules and supporting mechanisms. The a general structural design consisting of the hardware layout, the software technologies, and component allocation in the distributed environment is presented here as context for the translation rules.

5.1 *Hardware*

The hardware for the system usually involves multiple layers. EJB is generally an N-tier client server approach, although that has no relation to the hardware involved. All tiers may be running on one machine, or one tier on multiple machines.

5.2 *Software*

The business servers can be running the BEA Weblogic EJB Server or the J2EE supplied server. Supporting other servers is only a matter of generating the XML deployment descriptors.

The generated code is independent of any relational database product. Variations in the SQL support between vendors may require changes to the RDB schema generation templates.

6. MBSE to Relational Database Translation Rules

Relational databases define the persistent storage of data. The data stored is reflected on the Information Model for each domain.

6.1 *Attribute*

supertype: none; subtypes: none

An Attribute maps to a column in a class's table in a relational database. The data type of the attribute is mapped into the data types supported by the database. Attributes with the Identifier property are added to the constraints of the table.

Only attributes of persistent classes will be added to the schema. In addition, an attribute of a class may be marked with the boolean Transient property, in which case the attribute will not be included in the schema. By default, all attributes of a persistent class are persistent.

Data types, such as ServiceHandle and Handle must be mapped to a database supported type. One simple possibility is to serialize the Java object pointed to by the handle into a Blob attribute in the database. Handles to realized classes that are persistent can be serialized using Java serialization and then stored in a database in binary form. Handles to classes in other analyzed domains can also be serialized. If the handle is to an EntityBean, however, the handle must map into storage of the PrimaryKey of the EntityBean or into the EntityBean's handle (javax.ejb.Handle interface).

6.2 *Object (UML Class)*

supertype: none; subtypes: none

Each persistent object, marked by the Persistent property, results in a table in the database. The table name is based on the domain name and class name. Each table has an additional primary key integer attribute. This attribute is dependent on the unique ID generation capability available in most databases.

If one or more attributes are marked with the Identifier property, those attributes are included in an SQL UNIQUE CONSTRAINT clause within the schema. Note that this approach supports only one unique key based in addition to the uniquely generated integer key. If more UNIQUE constraints are required, an additional Unique property can be added, with a string value of the form: *attr1, attr2; attr1, attr3*, where the unique

attribute constraints are separated by “;” and the attributes within the constraint are separated by “,”.

6.3 Relationship

supertype: none; subtypes: BinaryRel, SubSuperRel

6.3.1 BinaryRel

supertype: Relationship; subtypes:none

Binary relations are formalized by placing the primary key attribute in the right table to support relational database Normal Form. For many to many relationships, an additional table will have to be created to store the relationship, if it does not already exist in the models.

6.3.2 SubSuperRel

supertype: Relationship; subtypes:none

Classes in Supertype/Subtype relationships are mapped into multiple tables, one table per class. The SubSuper relationship is implemented as a 1:1 unconditional relationship in the database schema. Each subtype has a formalizing attribute to its supertype class table. MBSE constraints apply to the instances in these tables – each subtype instance must have a corresponding supertype, and each supertype must have a corresponding instance in one of the subtype tables. Supertype tables add a column, SubType, that defines the table to look for the related subtype instance.

Traversing the relationship from super to subtype in the database requires two steps, as the supertype instance may be associated with only one of many instances in other tables. The first step is to extract the supertype data from the table, along with the table identifier of the subtype table. This is then used to access the subtype table for the associated instance.

An alternative that was discarded was collapsing supertypes into the subtype tables, duplicating the columns among subtypes. This removed the downcasting problem from super to subtype, but, since all relationships to the supertype must be duplicated downward, traversing from one table to the supertype gets the same complexity.

Classes that participate in SubSuper relationships must use bean managed persistence, due to the difficulties in configuring container managed persistence with inheritance.

6.4 Current State

supertype: none; subtypes: none

The current state of persistent, active objects is stored in the database as an integer.

6.5 InterDomain Relationships

Multiple domains will be stored in the database. Without interdomain relationships, each domain is an island, unconnected to the others. Specifying relationships between classes in different domains can speed up the access process by optimizing out the Java processing. This may not be applicable to all systems.

An example of an interdomain relationship is the relationship between a Customer in the CustomerProfile domain and an Order in OrderManagement. Separation by subject matter dictates that these are in a different domain, but there is a relationship between them nonetheless.

There are a couple of Analysis patterns that could be applied to interdomain relationships. First, an OrderingCustomer class could be added to the order management class. This is not the same as the Customer class in the CustomerProfile domain, since you cannot have the same class in multiple domains, as it violates subject matter separation. The OrderingCustomer is a reflection or view of the customer separate from his profile, and only captures aspects of the customer relating to the order. System requirements dictate that each OrderingCustomer must also be a Customer, so there is a 1:1c association from Customer to OrderingCustomer (a Customer that has never ordered anything is not an OrderingCustomer).

The second pattern creates a 1:Mc interdomain relationship directly between the Order and the Customer. An attribute is added to the order that is a handle to the Customer. In addition, a InterdomainRelationship property is added to the attribute to indicate that it is a handle to the CustomerProfile.Customer class. This property will formalize the Order-Customer relationship in the database by adding a column for the Customers key to the Order table. This allows standard SQL queries to be made directly on the database outside of analysis.

7. MBSE to Java/EJB Translation Rules

This section is broken down by the into sections based on the following MBSE constructs, Domain Chart, Information Modeling, State and Service Modeling, and Action Language. While this breakdown overlaps somewhat, it is easiest to consider these as sections.

This document only covers the large grained mappings for MBSE/UML to Java and EJB. Each section describes the mapping in a straight Java design, then the mapping using EJB. See the Springboard Users Guide for the full list of analysis elements accessible in the templates.

In cases where there are alternative translation rules that can apply, properties can be assigned to MBSE analysis constructs to guide the templates in the translation. Properties are name-value pairs associated with analysis elements, where the names are project-specific string constants. Properties can be specified in the model editor, or at translation time in the templates, using the SetProperty command in the Springboard Template Language. In each section, if there are possible alternative implementations, these are discussed, as well as the trade-offs, and property names are defined to guide the translation.

7.1 Domain Chart Constructs

7.1.1 Bridge

supertype: none; subtypes: none; A requirement flow line connecting two Domains on the domain chart.

A bridge defines the requirements level relationship between domains, as well as the service invocations that are made from one to the other. There is no implementation reflection of bridges in the implementation, except for package imports.

7.1.2 Constant

supertype:

System and domain wide constants are mapped to *public static final* member variables of the System of Domain classes. The constants are defined in the initialization hook .pal file along with the value.

7.1.3 DataType

supertype: none; subtypes: BaseType, GroupType, GroupIterType, InstanceReferenceType, ServiceHandle, UserDefinedType; The data type for an atomic data item

7.1.3.1 BaseType

supertype: DataType; subtypes: none; A built-in, predefined type

All base data types map to base data types in Java, int, String, double, etc.

Since Java passes everything by value, it is not possible to support output parameters from services with native Java data types. See the section on Services for more information.

7.1.4 UserDefinedType

supertype: DataType; subtypes: UserEnumerate, UserNonEnumerate; A data type defined by the user

7.1.4.1 UserEnumerate

supertype: UserDefinedType; subtypes: none; An enumerate defined by the user

Enumerated data types are not supported in Java. Instead, they are mapped to a series of *public static final int* member variables of either the System or Domain class, depending upon the scope of their definition.

Another possible implementation to consider here is to make each enumerated type its own Java class. This allows the language's type checking to verify that any use of an enumerate is valid within that context. The class itself contains factory methods to return the human-readable name of the enumerate, the next enumeration in the list, etc.

7.1.4.2 UserNonEnumerate

supertype: UserDefinedType; subtypes: none; A non-enumerate data type defined by the user

Since there is no concept of a typedef in Java, all non-enumerated user defined types are implemented as their corresponding Java base types.

All realized types will be handles to classes. The base type handle will be used for these types.

In the EJB implementation, the data type must be Serializable, or must be transient everywhere it is used.

7.1.5 Domain

supertype: DataTypeScope; subtypes: none

Domains are mapped into Java packages that contain the classes within the domain. In addition, a domain Java class is created that contains the domain services and the domain data types.

In EJB, the domain class will be mapped to a session bean if it has the EjbSessionBean property set to TRUE. The bean may be stateless or stateful, depending on results of further performance and throughput investigations. A property called EJBSessionBeanType could also be defined if the translation templates needed to support both session bean types.

Each EJB session bean that corresponds to a domain will have an instance of a PfdTask object to manage the event queue and timers for services that are invoked through it. The task will execute until there are no more events in the queue and no timers left active. Since the bean is stateless, the PfdTask instance will be refreshed to its initial state of an empty queue when finished with the service.

The domain's session will serve as a wrapper for the domain's implementation in a non-EJB class. The implementation class will look much like it does in the straight Java design, with static methods. Within a process, all services invoked in a domain local to the process will be through the implementation class. Across processors or processes, a session bean will be used. This will reduce the incidence of loopback/reentrance issues, though cross-process loopbacks are still possible.

For Example

```
class DomSessionBean extends SessionBean
{
public PfdTask task = new PfdTask();
void Service()
{ task.clear(); DomImpl.Service(); task.processOOA(); return;};
void Another () { ... };
}

class DomImpl
{
public Service()
{
// regular generates, etc
DomImpl.Another(); // local service

// For cross process, create session bean + invoke session service
...
}
public void Another() { ... };
}
```

OReilly EJB, p362 recommends avoiding chained stateful session beans. Chained session beans are beans that invoke other beans, all stateful. If a session times out, the whole state of that set of beans is corrupted and irrecoverable. Since domain invocations will result in chained session beans, the recommendation seems to be stateless sessions.

7.1.6 System

supertype: DataTypeScope; subtypes: none

The system is mapped to a package with the same name as the domain chart. This package contains all of the domain packages and the Sys class, which defines the system wide enumerated types and support for the mechanisms.

7.2 Information Modeling Constructs

7.2.1 Attribute

supertype: none; subtypes: none

An Attribute maps to a class attribute in Java. Currently the attribute is package visible and accessed directly in the implementation.

In EJB entity beans, attributes also map to a property or field in persistent classes. All attributes that are to be container managed must be declared public to allow the EJB container to manage them. Read and Write accessors will be implemented using the set/get<attrName> pattern in the bean interface. Since it is possible that some attributes may not be persistent, an optimization would be an Attribute level boolean Transient property.

If the IM is the basis for the RDB schema, using the relational database schema translation templates, then the mappings for EJB can be easily generated. If another schema is used, the EJB mapping becomes more complicated, and should be done using the EJB servers deployment capabilities.

Attributes of container managed entity beans are mapped to database tables through deployment descriptors, which are in XML format for EJB 1.1. Given straightforward mappings, we can generate most of this XML deployment descriptor directly.

7.2.2 Object (UML Class)

supertype: none; subtypes: none

A UML class is mapped to a Java class. All MBSE modeled classes are subtypes of PfdObject, for passive classes, or PfdActiveObject, for classes with state machines. PfdActiveObject also inherits from PfdObject.

In EJB, the PfdObject and PfdActiveObject are interfaces that are implemented by the EntityBeans.

A class is persistent if the Persistent property is set to TRUE. The translation rules will assume that a table has been created via the RDB templates to correspond to this class.

Persistent classes are mapped into EJB entity beans. Entity beans can be either container managed or bean managed. By default, the translation rules will use container managed persistence, since this is generally easier, and since there is a 1:1 mapping between the classes and the RDB schema. If container managed persistence is required, an EJBEntityType property will be added.

MBSE/UML classes mapped to entity beans will generate 4 classes/interfaces within the domain package:

Remote interface <classname>

Primary Key class <classname>PK

Home interface <classname>Home

Bean class <classname>Bean (Container or bean managed)

See documentation on EJB Entity beans for definitions and uses of these classes and interfaces.

UML classes that participate in sub/super relationships cannot be mapped to a container managed Entity bean. Based on research and prototyping, container management of the superclass bean data does not operate correctly. Bean managed persistence will be used for classes in a sub/super relationship in the analysis.

The remote interface for translated classes will need to integrate with the mechanisms. Interfaces equivalent to PfdObject and PfdActiveObject classes will have to be developed and added to the mechanisms. Attributes of these classes will be added during translation to generated bean classes, as will implementation of the interface methods.

7.2.3 Relationship

supertype: none; subtypes: BinaryRel, SubSuperRel

7.2.3.1 BinaryRel

supertype: Relationship; subtypes:none

By default, a binary association maps to a Java object reference. This assumes that all instances of all classes within the domain are in memory at the same time.

Since, in an EJB implementation, all instances related to an object may not be in memory at once, there must be a way to traverse an association from an in memory class to related classes in the database. To achieve this, formalizing attributes, already added to the tables through the RDB schema translation, will be added to the EJB classes. See the action language section on Relationships for more information on the use of formalizers to traverse relationships, link, and unlink.

References to transient classes will not have a formalizing attribute.

Associations that are [0|1]..* to [0|1]..* (many to many) do not require an associative object in MBSE. The associative table is created by the RDB templates, but no corresponding Java/EJB class will exist.

7.2.3.2 SubSuperRel

supertype: Relationship; subtypes:none

Subtype or inheritance relationships are mapped to the *extends* keyword in Java. Java does not support multiple inheritance, but then neither does MBSE/UML.

EJB Entity beans implement the EJB interfaces, and can therefore extend other modeled classes or PfdObject or PfdActiveObject.

As mentioned in the Object section, classes in sub/super relationships must be bean managed entity beans,

7.3 State and Service Modeling Constructs

The combination of domain services, state machines and actions, and class services make up the dynamic model of the domain. These constructs model the dynamics of the business logic of the particular domain. In a typical EJB design, all business logic is given to the session beans, keeping the entity beans very simple containers of data and relationships.

So, two basic implementation alternatives exist for mapping the state actions and instance base class services: attach them to the entity bean or create a separate class for them. Since the instance services and state actions cannot be called from outside the domain, there is no need for them to be session beans. The class consisting of state and instance

services will be an EJB client and an adjunct to the entity bean. This is in keeping with the EJB model of keeping entity beans as very simple data stores.

7.3.1 Action

supertype: none; subtypes: Service, State; the parent of all analysis elements with action language

7.3.1.1 Service

supertype: Action; subtypes: DomainService, ObjectService

7.3.1.1.1 DomainService

supertype: Service; subtypes: none

For the straight Java design, domain services map to static methods of the domain class. For realized domains, an interface with the same methods and signatures is defined. To realize the interface, this domain interface must be satisfied by a hand-coded class. At startup, an instance of this class is created and registered with the realized domain. The static methods invoked by the analyzed services are then passed to the registered class that implements the domain interface.

For the EJB design, domain services will be mapped to instance based methods of the domain's session bean class. Services that generate events will create a PfdTask instance to manage the event queue and wait until all events are processed and timers expired before returning.

Instance based services in EJB will be mapped to an adjunct class for entity beans. The adjunct class separates the business logic from the bean accessors.

7.3.1.1.2 ObjectService

supertype: Service; subtypes: none

In the Java design, object services will map to Java class methods, static methods for class based services.

In EJB, the class based services will be promoted to the domain level, but with only package visibility. Instance based services will remain with the entity bean or with its corresponding session bean.

Actually, class based services aren't visible outside of the domain, and can therefore be implemented within the second inner layer of the EJB interface design (the implementation layer).

7.3.1.2 State

supertype: Action; subtypes: none

In the Java design, state actions are implemented as instance based methods invoked by the state machine mechanisms.

State actions in EJB will be mapped to an adjunct class for entity beans. The adjunct class separates the business logic from the bean accessors.

7.3.1.3 Initialization Hooks

Initialization hooks are pieces of code that are configured (generated) to run automatically on startup. The hooks can be defined at the system or domain levels.

For stateless session beans, all defined initialization hooks must be run before the service executes.

7.3.2 Event

supertype: none; subtypes: none

Events are transient within a domain service invocation. All events will carry with them the reference to the PfdTask on which they are queued, so that subsequent event generates resulting from handling the event will be placed on the same queue.

Events destined for active objects that are also entity beans must carry the destination instance reference as an EJBHandle rather than a Java pointer. This is because the EJB container may choose to passivate (remove from memory) the target entity bean between generation and reception of the event.

7.3.3 Generate

supertype: EventAccessor; subtypes: none; an invocation of an event generate accessor

Event generation places events on the event queue within the context of a specific Task, shared among all the beans involved in a particular domain service. See the section on MBSE Semantics, subsection Task for more information on this.

7.3.4 State Transition Table

The table describes the state machine in terms of current state, incoming event, and next state.

This table will be stored statically in the active object itself. In the case of entity beans, the transition table will be stored in the same class as the state actions themselves.

7.3.5 Output Parameters

Since Java passes everything by value, it is not possible to support output parameters from services with native Java data types. Instead, a set of lightweight carrier classes are defined to wrap the base Java types and allow these classes to carry the output data back to the calling service. This restricts calls to services with output parameters to be standalone calls, not part of a more complex expression like an If condition.

In Java RMI and EJB, passing a carrier object will not be sufficient, as the changes are not propagated back to the caller. The only recourse to this is to pass the results in a wrapper class as the return value for the method in question. A serializable return vector must be used to transport output parameters and return values from server back to client.

A return vector of the service of java.util.Hashtable, using name-value pairs within the hashtable as the output parameters, could be used. This return vector will also carry the return value of the service and any service handles that were invoked on the server and destined for the client. The advantage of this interface is that it would be consistent across all client-server services, but the hashtable interface and name-value pairings suffer from possible run-time mismatches in naming if interfaces change.

Alternatively, a serializable class could be generated to carry the return vector. This is the current design plan. All output parameters would map to public member variables, as would the return value. For realized domains, an interface layer that simplifies the calling interface will be generated. See Realized Interfaces for more information.

A superclass mechanism, PfdReturnVector is defined in the SoftwareMechanisms. This is the superclass to all return vectors. It also carries the set of service handles invoked on the server and bound for the client.

7.4 Action Language

7.4.1 AttributeSelection

supertype: Expression; subtypes: none; an expression that reads or writes an attribute value

In both the Java and EJB designs, attribute reads and writes map to invocations of the attribute's get or set method.

7.4.2 Create

supertype: ObjectAccessor; subtypes: none; an invocation of an object create accessor

In the Java design, create statement maps to the creation of a new in- memory instance.

In EJB, the create statement of a class that is an entity bean maps to the ejbCreate invocation that creates the instance in the underlying database.

7.4.3 CreateServiceHandle

supertype: Statement; subtypes: none; an invocation of this built-in service

The CREATE ServiceHandle AL statement in both Java and EJB maps to the instantiation of a PfdServiceHandle, with the domain and service numbers filled in, along with any parameters.

See the discussion in Mechanisms, Service Handles for more information.

7.4.4 Delete

supertype: ObjectAccessor; subtypes: none; an invocation of an object delete accessor

In the Java design, delete statement maps to the removal of references to the in-memory instance, resulting in its availability for garbage collection.

In EJB, the delete statement of a class that is an entity bean maps to the remove invocation that deletes the instance in the underlying database.

7.4.5 Expression

supertype: none; subtypes: AttributeSelection, BinaryExpression, Constant, EventAccessor, LocalVariable, ObjectAccessor, ParameterVariable, RelationshipAccessor, ServiceInvocation, UnaryExpression; an invocation of a function and/or something that returns/has a value (as an rvalue), or something that can have its value set (as an lvalue)

Expression subtypes are mapped into Java.

7.4.6 Instance Lookups - Find/Foreach

supertype: ObjectAccessor; subtypes: none; an invocation of an object find accessor

A lookup statement in MBSE may have two sources, a relationship navigation or the class instances, and each source Find may be of two types, a find first/last (FIND), find all (FOREACH), and each Find type may include an optional Where clause for further limitations on the returned instances.

In the straight Java design, lookup statements map to searches on the in memory relationship or instance lists. The instances are iteratively compared with the Where clause.

In the EJB design, for entity beans, not all the class instances or related instances may be in memory, so lookup must map to database searches, using the EJB as an interface. Each type of lookup (class based find and relationship traversal) referenced in the AL of the domain will have an interface in the Home interface as well as the bean class.

Custom FINDS can be implemented as separate methods and use SQL to access the database directly. Matching entities must result in the return of a Collection of Primary Key objects, which the container would translate to instances. If required, we can map class based finds and association traversals directly to SQL.

7.4.7 InvokeServiceHandle

supertype: Statement; subtypes: none; a statement that invokes a ServiceHandle

The CALL statement invokes a service handle. If the service referred to by the handle resides in the same process, the handle is decoded and invoked. If the target service resides on an EJB server, a session bean is obtained at the service handle invoked. If the target service is invoked on the EJB server and resides on the EJB client, the service handle is added to the return vector of the service for transport back to the client.

See the discussion in Mechanisms, Service Handles for more information.

7.4.8 Link

supertype: RelationshipAccessor; subtypes: none; an invocation of a relationship link accessor

In Java, a link results in bi-directional pointers set in the participating objects (and in an associative object, if necessary).

In EJB, a link does the same thing, since both objects must be in memory to be linked. However, for entity beans, the LINK also sets the formalizing database attribute for the relationship. Links may also result in the creation of associative objects in the database to support many to many relationships.

7.4.9 Navigation

supertype: RelationshipAccessor; subtypes: none; an invocation of a relationship navigation accessor

In Java, relationship navigation, from within a FIND or FOREACH statement, is taken from the pointer list within the class.

In EJB, for entity beans and relationships between entity beans, the navigation is mapped to a find/lookup method in the Home interface and bean class to find the associated instances from the database.

7.4.10 ServiceInvocation

supertype: Expression; subtypes: none; the invocation of an object or domain service

An invocation of a class method of the domain or UML class, in Java.

In EJB, maps to the appropriate invocation of a service, either in a session bean or an entity bean. In some cases, it may need to create a new session bean before invoking the method.

7.4.11 SubSuperNavigation

supertype: Expression; subtypes: none; a "cast" from a supertype to one of its subtypes

Results in a simple downcast in both Java and EJB designs. NOTE: For EJB 1.1, this is done using the `javax.rmi.PortableRemoteObject.narrow()` method. See O'Reilly pp 133-135 for exceptions.

7.4.12 Unlink

supertype: RelationshipAccessor; subtypes: none; an invocation of a relationship unlink accessor

The reverse of LINK, in Java, it removes the bi-directional references between the instances.

In EJB, it also removes the formalizing attributes from the database table. Unlink will also remove any associative objects in the database.

7.5 MBSE Semantics

7.5.1 Timers

A single timer thread will be shared among all the session beans implementing the domain interfaces. The timer thread will accept timer settings and respond to the requesting Task when the timer has expired. This thread should be started outside the EJB environment, but shareable within it.

7.5.2 Instance Lists

An instance list is the set of all instances of a particular class. In this design, the set is only completely visible within the relational database that supports the system. As such, there are really 2 instance lists – the complete set in the database, and the partial set in memory. It is assumed that the EJB container takes care of insuring that an instance (entity bean) is not in memory twice.

Transient class instances are held completely in memory, but should not overlap between tasks or instances of the domain.

No in-memory instance lists are kept for entity beans, since the database keeps the list.

In memory lists for non-entity beans must be kept on a thread-by thread basis. The instance list must then be part of the thread's MBSE execution context.

7.5.3 Service Handles

EJB does not provide any support for asynchronous communication between two entities. The model is strictly client server, where all processing is synchronous and controlled by the client, invoking the server and blocking until processing is complete and the results returned. Thus, the concept of service handles does not map well to the communication between domains on the client and domains on the server.

Service handles invoked on the EJB server that are services of the client will be serialized across the network, carried in the return vector for the service. Once they arrive at the client, they are deserialized and invoked. Note that this limits a server thread from communicating with other clients, only the client that started the particular thread.

Domain services of the client will be wrapped in an interface class that will create the service handle and insert it into the return vector. As part of the return vector from the service, a set of zero or more service handles will be returned. The service handles will be serialized to cross the wire. Any service handles bound for the client will be returned in the return vector, unpacked, and invoked before the method returns on the client. The service would also package up the return values to return them to the client.

Between domains on the same side of the client-server boundary, the mapping is straightforward. Either the service handle is decoded and invoked directly or a session bean is created and then the service invoked.

7.5.4 Task

The task consists of the MBSE event loop and the interface to timers. Within the context of the EJB design, there could be many tasks running at the same time, coordinated through the EJB container and the underlying relational database.

Each domain session bean service accepts parameters, does processing, and, if it generates an event, creates a task, deposits the event, and spins the event loop. The event loop executes until there are no events left to process and there are no timers active.

Since there are multiple tasks running at the same time, the objects must know what task to place new events into. There is no global area or global key that can be used to store the PfdTask reference, so it must be passed around during processing. The PfdTask reference is included with each event, so the receiving states will know what PfdTask to enqueue the next event on. Any services that can generate an event will have a PfdTask added to the interface. Session bean interfaces do not expose the PfdTask, but include their PfdTask in calls to domain implementation services.

We may need to add the PfdTask to all services, since a service may call another service that may generate an event. Or, the whole tree of processing for a service may need to be explored during translation.

7.6 Java Specific Issues

7.6.1 Exceptions

Exception handling is not part of the MBSE semantics. Java exceptions are not thrown or caught by the mechanisms or any of the design templates.

7.7 EJB Specific Issues

7.7.1 Loopbacks

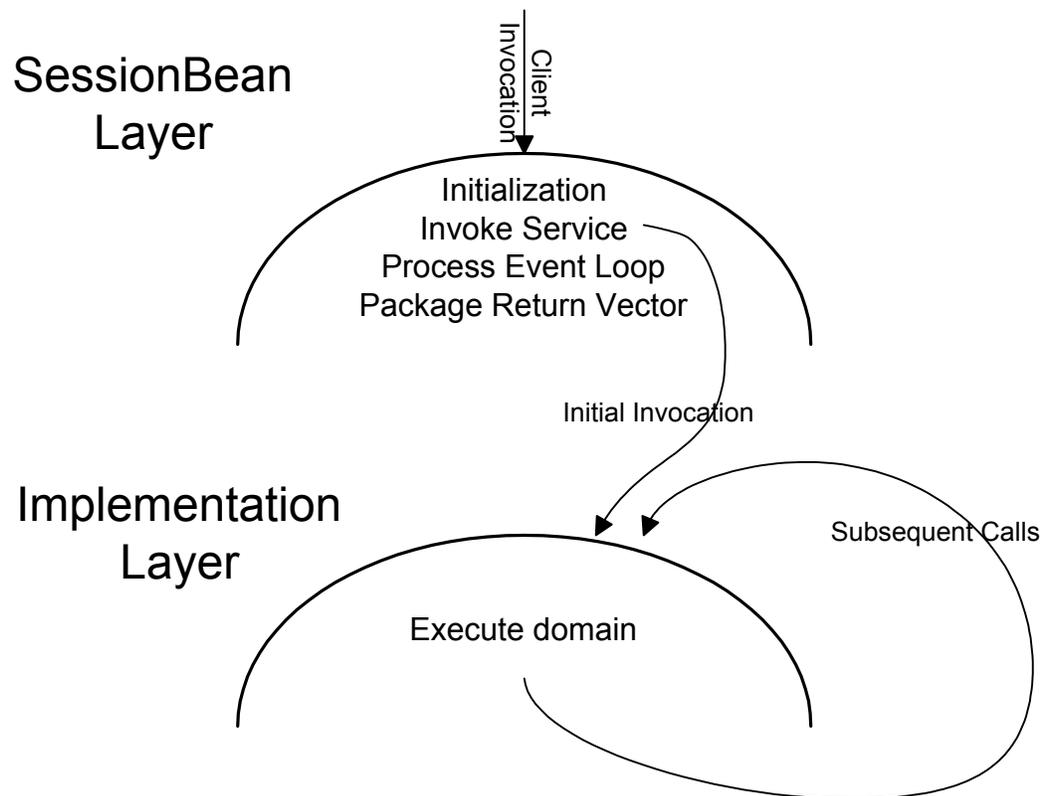
In EJB, the container manages all the resources for the application, memory, threading, transactions, etc. As a result, this places some constraints on the application. One of these constraints prohibits loopbacks in the EJB processing. Loopbacks are when instance A calls a service of instance B, which then calls another service of instance A. Loopback actions are not encouraged, as they damage the threading and synchronization model of the EJB container. Settings

on the bean deployment can allow loopbacks, but this removes much of the protection EJB provides.

To avoid loopbacks, the domain API is separated into 2 levels. Clients of the EJB server, mainly the user interface domains, use the Session Bean interface. The SessionBean provides the context, threading, and transactions for the operations. The second level is the actual domain implementation.

The SessionBean simply invokes the domain implementation level, then executes the PfdTask event loop. When the event loop is complete, it packages the return vector and passes it back to the client.

All processing within the domain uses the second level domain interface, avoiding the session bean loopback problems within the domain. If multiple session beans are used, then they should be stateless and each create a new session bean to converse with the other domains to avoid loopbacks.



7.7.2 Exceptions

Exception handling is not part of the MBSE semantics. No application exceptions are not thrown or caught by the mechanisms or any of the design templates. EJB exceptions are caught and rethrown as required by the EJB specification.

7.7.3 Transactions

The EJB container and server provide distributed transaction context for the application. The EJB container provides a set of properties on EJB classes and methods that can be used to control and optimize access to the database while providing enough protection to insure that data remains consistent. Transaction membership and isolation control are the properties that can be set. The level of support for these varies between vendors and must be investigated for TopLink and Oracle.

These EJB properties are set in the deployment descriptor for the bean. A set of properties can be captured in the MBSE models that correspond to the EJB properties. The deployment descriptor can be generated from these.

We assume that helper classes – non EJB classes called from a bean - carry transaction processing along, and that the transaction will propagate to other beans called by the helper classes as per the EJB transaction propagation rules and settings.

Deadlocks are detected by the EJB container or the underlying database when tested with the BEA WebLogic server and the Cloudscape database. In a deadlock situation, one of the session clients will time out, causing the server to roll back any pending transactions.

For more information on transactions and EJB, see chapter 8 of [Enterprise JavaBeans](#).

7.8 Realized Interfaces

The interface between client and service in mapping MBSE to EJB gets a bit complicated with service handles and output parameters, neither of which are directly supported by EJB. This is not a problem for analyzed domains, since the code can be generated to deal with that. With realized domains, such as the user interface, it becomes more of a problem. To ease the realized to analyzed interface across the EJB client-server boundary, a simplified wrapper class around the session bean will be provided.

The adapter class will have the same interface as the MBSE models, with carrier classes for output parameters. The decoding of the return vector and invocation of any service handles will be performed before the service returns.

8. MODELING CONVENTIONS AND RESTRICTIONS

Springboard expects your OOA models follow the conventions described in "Pathfinder Solutions OOA/UML Modeling Guide".

Ideally, there are no conventions and restrictions placed on the analysis by the implementation or design technology. Practically, there is a small amount of feedback, especially for a design that is under development and intended to be used before it is complete. The following are conventions and constraints on the model for use with the Java EJB design. This list will change as development progresses.

- Use of output parameters from services in complex expressions. Service calls with output parameters cannot be used within a compound expression, like an IF or WHILE statement.
- Loopback limitation on entity beans and session beans. Make as much inter-class communication as possible asynchronous via events to avoid this.