

IBM Cúram Social Program Management  
Version 7.0.0

*Cúram Modeling Reference Guide*

**IBM**

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 91

**Edition**

This edition applies to IBM Cúram Social Program Management v7.0.0 and to all subsequent releases unless otherwise indicated in new editions.

Licensed Materials - Property of IBM.

© **Copyright IBM Corporation 2012, 2016.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

© Cúram Software Limited. 2011. All rights reserved.

# Contents

<b>Figures</b> . . . . .	<b>vii</b>
--------------------------	------------

<b>Tables</b> . . . . .	<b>ix</b>
-------------------------	-----------

## Cúram Modeling Reference . . . . . 1

Cúram Modeling Overview . . . . .	1
UML Overview . . . . .	1
UML and the Input Meta-model . . . . .	2
Overview of the Architecture Layers . . . . .	2
Remote Interface Layer . . . . .	2
Business Object Layer . . . . .	2
Data Access Layer . . . . .	3
Stereotypes . . . . .	3
Class stereotypes . . . . .	3
Attribute Stereotypes . . . . .	4
Operation Stereotypes . . . . .	4
Relationship Stereotypes . . . . .	5
Data types . . . . .	6
Packages. . . . .	8
Options . . . . .	8
CODE_PACKAGE . . . . .	8
Audit Mappings Classes . . . . .	10
Rules . . . . .	10
Outputs . . . . .	11
Options. . . . .	11
Domain Definition Classes . . . . .	12
Options. . . . .	12
Code Table Name . . . . .	12
Code Table Root. . . . .	13
Compress Embedded Spaces. . . . .	13
Convert to Uppercase . . . . .	13
Custom Validation Function Name . . . . .	14
Default . . . . .	14
Maximum Size . . . . .	14
Maximum Value. . . . .	14
Minimum Size . . . . .	14
Minimum Value . . . . .	14
Multibyte Expansion Factor . . . . .	14
Pattern Match . . . . .	15
Remove Leading Spaces . . . . .	15
Remove Trailing Spaces . . . . .	15
Storage Type . . . . .	15
Overriding a Domain Definition . . . . .	15
How to use Domain Definition Overrides . . . . .	15
Considerations / Limitations . . . . .	15
Usage Rules . . . . .	16
Entity Classes . . . . .	16
Rules . . . . .	16
Attributes . . . . .	16
Details . . . . .	17
Key . . . . .	17
Operations . . . . .	17
Database Operations . . . . .	17
Non-database Operations. . . . .	17
Outputs . . . . .	17

Standard Key Structs . . . . .	17
Standard Details Structs . . . . .	17
Standard List Structs . . . . .	18
Options. . . . .	18
Abstract . . . . .	18
Allow Optimistic Locking . . . . .	18
Audit Fields . . . . .	18
Enable Validation . . . . .	18
Last Updated Field . . . . .	18
No Generated SQL . . . . .	19
Replace Superclass . . . . .	19
Concurrency Control - Optimistic Locking . . . . .	19
Table Level Auditing . . . . .	20
Information Captured by Table-level Auditing . . . . .	20
Storage of Audit Information . . . . .	21
Exit Points. . . . .	21
Pre Data Access . . . . .	21
Post Data Access . . . . .	21
Validation . . . . .	21
On-fail . . . . .	22
Exit Point Parameters . . . . .	22
What should exit points be used for . . . . .	22
What should exit points <i>not</i> be used for . . . . .	22
Entity Inheritance . . . . .	22
Rules when Using Entity Inheritance . . . . .	23
Last Updated Field . . . . .	23
Extension Classes . . . . .	24
How to use Extension Classes . . . . .	24
When to use Extension Classes . . . . .	24
Considerations / Limitations . . . . .	24
Usage Rules . . . . .	25
Facade Classes . . . . .	25
Rules . . . . .	25
Operations . . . . .	25
default . . . . .	25
batch . . . . .	25
wmdpactivity. . . . .	25
qconnector. . . . .	25
Options. . . . .	26
Abstract . . . . .	26
Generate Facade Bean . . . . .	26
Replace Superclass . . . . .	26
Process Classes . . . . .	26
Business Process Objects . . . . .	26
Rules . . . . .	26
Operations . . . . .	27
default . . . . .	27
batch . . . . .	27
wmdpactivity. . . . .	27
qconnector. . . . .	27
Options. . . . .	27
Abstract . . . . .	27
Generate FIDs . . . . .	27
Replace Superclass . . . . .	27
Struct Classes. . . . .	27
Rules . . . . .	28

Outputs . . . . .	28	Parameter and Generator Notes . . . . .	43
Options . . . . .	28	Standard Readmulti . . . . .	44
Audit Fields . . . . .	28	Description . . . . .	44
Attributes . . . . .	28	Use . . . . .	44
Attribute Rules . . . . .	29	Parameter and Generator Notes . . . . .	44
Attribute Options . . . . .	29	Non-standard Read (Generated SQL) . . . . .	44
Allow NULLs . . . . .	29	Description . . . . .	44
Multibyte Expansion Factor . . . . .	30	Use . . . . .	44
Operations . . . . .	30	Parameter and Generator Notes . . . . .	45
Rules . . . . .	30	Non-standard Readmulti (Generated SQL) . . . . .	45
Operation Options . . . . .	30	Description . . . . .	45
Audit BI (Business Interface) Calls to this		Use . . . . .	45
Operation . . . . .	30	Parameter and Generator Notes . . . . .	45
Auto ID Field . . . . .	31	Non-key Read . . . . .	46
Auto ID Key . . . . .	31	Description . . . . .	46
Business Date . . . . .	31	Use . . . . .	46
BytesMessage encoding character set . . . . .	32	Parameter and Generator Notes . . . . .	46
Database Table-level Auditing . . . . .	33	Non-key Readmulti . . . . .	46
Field Level Security . . . . .	34	Description . . . . .	46
JNDI name of the QueueConnectionFactory		Use . . . . .	47
class . . . . .	34	Parameter and Generator Notes . . . . .	47
JNDI name of the transmission queue . . . . .	34	Entity Update Operations . . . . .	47
JNDI name of the reply queue . . . . .	35	Standard Modify . . . . .	47
Message type . . . . .	35	Description . . . . .	47
No Generated SQL . . . . .	35	Use . . . . .	47
On Fail Operation . . . . .	36	Parameter and Generator Notes . . . . .	48
Optimistic Locking . . . . .	36	Non-standard Modify (Generated SQL) . . . . .	48
Order By . . . . .	36	Description . . . . .	48
Post Data Access Operation . . . . .	37	Use . . . . .	48
Pre Data Access Operation . . . . .	37	Parameter and Generator Notes . . . . .	48
Readmulti_Max . . . . .	37	Non-key Modify . . . . .	48
Readmulti_Informational . . . . .	37	Description . . . . .	48
Response message timeout (seconds) . . . . .	37	Use . . . . .	49
Security . . . . .	37	Parameter and Generator Notes . . . . .	49
SQL . . . . .	38	Entity Delete Operations . . . . .	49
Transactional . . . . .	38	Standard Remove . . . . .	49
Where . . . . .	38	Description . . . . .	49
Operation Parameter Options . . . . .	38	Use . . . . .	49
Mandatory Fields . . . . .	38	Parameter and Generator Notes . . . . .	49
Entity Operations Overview . . . . .	39	Non-standard Remove (Generated SQL) . . . . .	50
Standard Operations . . . . .	40	Description . . . . .	50
Standard Single-Record Operations . . . . .	40	Use . . . . .	50
Standard Multi-Record Operations . . . . .	40	Parameter and Generator Notes . . . . .	50
Non-Standard Operations . . . . .	40	Non-key Remove . . . . .	50
Generated SQL Operations . . . . .	40	Description . . . . .	50
Handcrafted SQL Operations . . . . .	41	Use . . . . .	51
Non-Key Operations . . . . .	41	Parameter and Generator Notes . . . . .	51
Batch Operations . . . . .	41	Entity Batch Operations . . . . .	51
Entity Insert Operations . . . . .	41	BatchInsert . . . . .	51
Standard Insert . . . . .	41	Description . . . . .	51
Description . . . . .	42	Use . . . . .	52
Use . . . . .	42	Parameter and Generator Notes . . . . .	52
Parameter and Generator Notes . . . . .	42	BatchModify . . . . .	52
Non-standard Insert (Generated SQL) . . . . .	42	Description . . . . .	52
Description . . . . .	42	Use . . . . .	52
Use . . . . .	42	Parameter and Generator Notes . . . . .	52
Parameter and Generator Notes . . . . .	43	Entity Handcrafted SQL Operations . . . . .	53
Entity Read Operations . . . . .	43	Non-standard . . . . .	53
Standard Read . . . . .	43	Description . . . . .	53
Description . . . . .	43	Use . . . . .	54
Use . . . . .	43	Parameter and Generator Notes . . . . .	54

Non-standard multi . . . . .	54	The Meaning of super . . . . .	76
Description . . . . .	54	Enforcing the Factory Mechanism . . . . .	76
Use . . . . .	54	Summary . . . . .	76
Parameter and Generator Notes . . . . .	55	Cúram JMS Queue Connectors . . . . .	76
Example 1 - nsmulti with a Single (List)		Domain Definition Overview . . . . .	77
Parameter . . . . .	55	How It Works / What It Does . . . . .	77
Example 2 - nsmulti with Two Parameters		Options on qconnector Operations. . . . .	77
(Key + List) . . . . .	57	How to Use qconnector Operations . . . . .	78
Using Handcrafted SQL in Non-Standard Entity		Decide on Format of Message and Create the	
Operations . . . . .	59	struct(s) to Correspond to the Message . . . . .	78
Overview . . . . .	59	Add the operation to the application	
Using Host Variables . . . . .	59	meta-model. . . . .	79
“Null” Considerations. . . . .	59	Configure the Queues in the Application	
For Update Considerations With DB2 for		Server . . . . .	79
z/OS . . . . .	60	Implement the message recipient in the	
SQL Example 1 . . . . .	60	remote system . . . . .	79
SQL Example 2 . . . . .	61	Rules / Restrictions . . . . .	79
Aggregation . . . . .	62	Encoding Methods for Fundamental Types . . . . .	79
A Special Case . . . . .	62	Using Customized Encoding/Decoding Classes . . . . .	80
One-to-One Aggregation . . . . .	62	Example 1 - Working with Variable Length Fields . . . . .	81
One-to-Many Aggregation . . . . .	64	Example 2 - Working with Lists . . . . .	84
Assignable. . . . .	65	Subclassing . . . . .	87
Explicit Field Assignment. . . . .	66	How to Model It . . . . .	87
Suppressing Default Assignment Fields . . . . .	68	Basic Subclassing . . . . .	87
Combining structs . . . . .	68	Replacing the Superclass . . . . .	87
Foreign Keys . . . . .	69	Abstract Classes . . . . .	88
How to Add a Foreign Key to a Database Table . . . . .	70	Restrictions . . . . .	88
Naming Primary and Foreign Key Constraints. . . . .	70	How to write Code for Subclassing . . . . .	88
Example . . . . .	70	Example - Using Subclassing to Override Entity	
Indexes. . . . .	71	Exit Points. . . . .	88
How to Add an Index to a Database Table . . . . .	71	Overriding Validation Exit Point . . . . .	88
Naming Indexes. . . . .	71	Overriding Pre Data Access, Post Data Access,	
Example . . . . .	72	and On-Fail exit points . . . . .	88
Unique Indexes . . . . .	72	Application Customization . . . . .	89
Generated Class Hierarchy . . . . .	73		
Basic Hierarchy Example . . . . .	73	<b>Notices . . . . .</b>	<b>91</b>
Hierarchy for Subclasses . . . . .	75	Privacy Policy considerations . . . . .	92
Hierarchy for Abstract Classes . . . . .	75	Trademarks . . . . .	93
Considerations . . . . .	76		
Access Control - private/protected/public/			
package . . . . .	76		



---

## Figures

1.	extra.generator.options property in Bootstrap.properties . . . . .	23	16.	Struct for return result . . . . .	61
2.	Sample datamanager_config.xml for adding field level security information to the database.	34	17.	Java Interface . . . . .	61
3.	Inserting field level security SIDs into the infrastructure SecurityIdentifier table . . . . .	34	18.	SQL Implementation . . . . .	61
4.	Handcrafted data access implementation for a standard read . . . . .	35	19.	Java Interface . . . . .	61
5.	Handcrafted data access implementation for a readmulti . . . . .	36	20.	Struct for employer key . . . . .	62
6.	Operation Signature. . . . .	38	21.	SQL Implementation . . . . .	62
7.	Pseudo-Code for Parameter Structures . . . . .	39	22.	SQL Implementation with qualified parameters	62
8.	SQL for nsmulti with a single (list) parameter	55	23.	Example Java code for combining structs	69
9.	Pseudocode for generated structs for use by nsmulti operation . . . . .	56	24.	Equivalent Java code for combining structs	69
10.	Generated Java interface for nsmulti operation	56	25.	Using a factory to create an instance of MyClass . . . . .	75
11.	Calling a nsmulti operation from handcrafted Java code (one parameter). . . . .	56	26.	Sample QueueConnectorFieldMappers.properties . . . . .	81
12.	SQL for nsmulti with a key and list parameters	57	27.	Pseudo code for the struct to be mapped:	81
13.	Pseudocode for generated structs for use by nsmulti with key and list parameters . . . . .	58	28.	Pseudo code for the BPO interface . . . . .	81
14.	Generated Java interface for nsmulti operation with key and list parameters. . . . .	58	29.	The property file entries linking the fields to the mapper . . . . .	81
15.	Calling a nsmulti operation from handcrafted Java code (two parameters) . . . . .	59	30.	Mapper class implementation for variable string . . . . .	83
			31.	Pseudo code for the structs to be mapped:	84
			32.	Pseudo code for the BPO interface . . . . .	84
			33.	The property file entry linking the fields to the mapper . . . . .	85
			34.	Mapper class implementation for list of structs	86



---

## Tables

1. Cúram class stereotypes . . . . .	3	6. Domain primitive types at different levels of a	
2. Cúram Attribute Stereotypes . . . . .	4	Cúram application . . . . .	12
3. Cúram Operation Stereotypes . . . . .	4	7. Mapping of Class and Attribute Stereotypes	29
4. Cúram Relationship Stereotypes . . . . .	5	8. Data types and nulls . . . . .	60
5. Cúram Data Types . . . . .	6	9. Encoding methods . . . . .	79



---

# Cúram Modeling Reference

Use this information to learn about the Cúram UML meta-model. The Cúram generator uses this UML meta-model to automatically generate all the required stubs, skeletons, classes, and communications that are required to interact with a database and remote clients. This approach allows the developer to concentrate on providing the application business logic.

---

## Cúram Modeling Overview

You can create client/server applications with a minimum of complexity by using the Cúram UML model. The UML model simplifies database access, EJB management, and client/server interaction.

The Cúram Generator uses the UML meta-model to automatically generate all the required stubs, skeletons, classes, and communications that are required to interact with a database and remote clients. The Cúram Generator allows the developer to concentrate on the application business logic. This document describes the tools and components that are used to model the application, and how the Cúram Server Code Generator treats each component when it generates classes.

The UML meta-model is a platform-independent model that describes the following aspects of the application:

- Domains: application-specific data types. Analogous to C++ typedefs.
- Entities: The objects that are modeled and persistently stored by the application. Corresponds to relational database tables.
- Processes: related sets of activities to achieve some business goal.
- Structs: passed as messages throughout the application. Analogous to structs in C++.
- Remote Interfaces: client-visible interfaces through which server functions can be accessed.

You edit the UML meta-model with IBM®Rational® Software Architect. Code is generated by the Cúram Server Code Generator, which produces the following outputs:

- JavaServer implementation code
- Java beans
- XML for the database entities and other classes in the model.

The XML produced by the Cúram Server Generator is processed by the Cúram Data Manager. The Cúram Data Manager produces the relevant SQL scripts that are used to create the required database structure for the application.

This document also provides a reference for Cúram model-based functions such as: Cúram domains, classes, operations, attributes and how they map to the underlying database.

---

## UML Overview

UML constructs are created and maintained with Rational Software Architect. This model is referred to as the input meta-model.

The input meta-model is used as input to the Cúram generator. It is a logical representation of the system. The meta-model is the mechanism that developers use to tell the generator what to generate.

Connection

Connection

The meta-model consists of a set of packages, which contain class representations. Classes have attributes and operations that have relationships with one another. Classes in the input meta-model result in various generated Java classes and in some cases tables and indexes in generated DDL (Database Definition Language). DDL is an SQL language subset that enables the structure and instances of a database to be defined in a human and machine-readable form.

## UML and the Input Meta-model

UML constructs, created and maintained by the user with Rational Software Architect, are referred to collectively as the *input meta-model* and this is used as input to the Cúram generator. It is a logical representation of the system being developed. Another way of looking at it is: it is the mechanism that developers use in order to tell the generator what to generate.

This meta-model consists of a set of packages, which in turn contain *class* representations, potentially containing *attributes* and *operations*, which have relationships with one another. Classes in the input meta-model result in various generated Java classes and in some cases tables and indices in generated DDL<sup>1</sup>.

## Overview of the Architecture Layers

The Cúram architecture is conceptually divided into three layers, as follows.

### Remote Interface Layer

The Remote Interface Layer presents an interface to business functions that can be used by a client program. It also interacts with third-party middleware components to ensure consistency and atomicity of the transactions that run in business functions.

### Business Object Layer

The Business Object Layer implements all of the server's business functionality. As such, this layer contains the business application's "smarts". Within this layer, Business Process Objects (BPOs) represent the basic business entities modeled by the server application. BPOs implement the business logic of a Cúram server application. Typically BPOs are responsible for manipulating Entity Objects in a business-specific way. The Business Object Layer is where most of the development effort should be concentrated in business application development.

#### Related reference:

"Process Classes" on page 26

Process classes are a collection of operations that encapsulate a business process. Process classes do not have data maintenance operations, or any relationship with database tables. They manipulate other entity and process classes to implement a business process.

---

1. DDL means Database Definition Language. It is an SQL language subset enabling the structure and instances of a database to be defined in a human and machine-readable form.

## Data Access Layer

The Data Access Layer is responsible for all interactions with the back end Relational Database Management System (RDBMS).

### Related reference:

“Entity Classes” on page 16

An entity is a collection of fields and associated database operations. Entity classes are the fundamental building blocks of systems that are developed with Cúram. They correspond to database tables. The Cúram generator supports automatic code generation for entity classes.

## Stereotypes

Stereotypes are a Unified Modeling Language (UML) concept that is used to further describe the various aspects of the Cúram application model. In UML, a stereotype is a string expression that is used to assign a classification to an object.

In general, stereotypes affect the behavior of the generator and thus determine its output. Therefore, for example, an entity class is identified by having a stereotype of entity and DDL and data-access code produced by the generator.

The supported stereotypes are as shown in the following sections.

### Class stereotypes

The following table lists the class stereotypes with a short description and reference to where they are described in more detail.

Table 1. Cúram class stereotypes

Class stereotype	Description	Reference
audit_mappings	An audit mappings class enables more fields to be defined in the database for auditing purposes.	“Audit Mappings Classes” on page 10
domain_definition	A domain definition is a meta-model class that defines a data type.	“Domain Definition Classes” on page 12
entity	An entity class encapsulates data-maintenance functions on a database table.	“Entity Classes” on page 16
extension	Extension classes are intended to be used to change the Audit Fields or Last Updated Field options of an entity or struct class.	“Extension Classes” on page 24
facade	Facade classes are used to create client-visible operations. They provide a simplified interface to a larger body of code, such as a class.	“Facade Classes” on page 25
process	A process class encapsulates a business process.	“Process Classes” on page 26

Table 1. *Cúram class stereotypes (continued)*

Class stereotype	Description	Reference
rest	The rest stereotype class is for facades that contain nested lists that are used by REST API services.	Cúram REST API methods
struct	A struct class is a meta-model representation of a Java class that contains a collection of fields.	"Struct Classes" on page 27
webservice	A Webservice class represents an inbound legacy web service.	<i>Cúram Web Services Guide</i>
wsinbound	A WS Inbound class represents an inbound web service.	<i>Cúram Web Services Guide</i>

## Attribute Stereotypes

The following table lists the attribute stereotypes with a short description and reference to where they are described in more detail.

Table 2. *Cúram Attribute Stereotypes*

Stereotype	Description	Reference
audit_mappings	An audit field entry on the Audit Mappings Class.	"Audit Mappings Classes" on page 10
dataitem	An attribute on an RDO or ListRDO class.	Data item attribute
default	A public attribute or field in a struct class.	"Struct Classes" on page 27
details	An attribute or field that is part of an entity but not part of the entity key.	"Entity Classes" on page 16
key	An attribute or field that is part of an entity's key.	"Entity Classes" on page 16

## Operation Stereotypes

The following table lists the operation stereotypes with a short description and reference to where they are described in more detail.

Table 3. *Cúram Operation Stereotypes*

Stereotype	Description	Reference
batch	Process by the Batch Launcher by generated wrapper code	"batch" on page 25 and "batch" on page 27
batchinsert	For inserting large amounts of data by batch	"BatchInsert" on page 51
batchmodify	For modifying large amounts of data by batch	"BatchModify" on page 52
default	Standard non-database operation	"default" on page 25 and "default" on page 27
insert	Standard database insert	"Standard Insert" on page 41

Table 3. *Cúram Operation Stereotypes (continued)*

Stereotype	Description	Reference
modify	Standard database update	"Standard Modify" on page 47
nkmodify	Non-key database update	"Non-key Modify" on page 48
nkread	Non-key database read	"Non-key Read" on page 46
nkreadmulti	Non-key database read	"Non-key Readmulti" on page 46
nkremove	Non-key database delete	"Non-key Remove" on page 50
ns	Database operation for handcrafted SQL	"Non-standard" on page 53
nsinsert	Non-standard database insert	"Non-standard Insert (Generated SQL)" on page 42
nsmodify	Non-standard database update	"Non-standard Modify (Generated SQL)" on page 48
nsmulti	Database operation for handcrafted SQL	"Non-standard multi" on page 54
nsread	Non-standard database read	"Non-standard Read (Generated SQL)" on page 44
nsreadmulti	Non-standard database read	"Non-standard Readmulti (Generated SQL)" on page 45
nsremove	Non-standard database delete	"Non-standard Remove (Generated SQL)" on page 50
qconnector	For connecting to external JMS	"qconnector" on page 25 & "qconnector" on page 27
read	Standard database read	"Standard Read" on page 43
readmulti	Standard database read	"Standard Readmulti" on page 44
remove	Standard database delete	"Standard Remove" on page 49
wmdpactivity	Deferred processing	"wmdpactivity" on page 25 and "wmdpactivity" on page 27

## Relationship Stereotypes

The following table lists the attribute stereotypes with a short description and reference to where they are described in more detail.

Table 4. *Cúram Relationship Stereotypes*

Stereotype	Description	Reference
aggregation	The ability to embed or nest instances of one type of class within another type of class	"Aggregation" on page 62
assignable	Use an assignable relationship to map differing or exclude fields for an assign function.	"Assignable" on page 65
extension	The link between an extension class and target class.	"Extension Classes" on page 24
foreignkey	A modeled description of a database foreign key.	"Foreign Keys" on page 69

Table 4. *Cúram Relationship Stereotypes (continued)*

Stereotype	Description	Reference
index	A modeled description of a database index.	"Indexes" on page 71
uniqueindex	A modeled description of a database unique index.	"Unique Indexes" on page 72

## Data types

The input meta-model supports a number of data types that provide abstraction for the developer from the different underlying data types used by the database, middleware, and Java layers. These data types can be used to define attributes, arguments, and return values in a platform-neutral and database-neutral way. The SDEJ maps the attributes, arguments, and return values to the appropriate data type in each layer of the application.

Table 5. *Cúram Data Types*

Type	Description
SVR_BLOB	Used for holding binary data.  Corresponds to class <code>curam.util.type.Blob</code> .  Requires a size qualifier that is used only if the field is used on a database table.  Fields of type SVR_BLOB might be null on the database.
SVR_BOOLEAN	Used for holding binary values.  Corresponds to the primitive Java type <code>boolean</code> .  Is stored as a single character field on the database where 0 = false and 1 = true.  Fields of type SVR_BOOLEAN cannot be null on the database.
SVR_CHAR	Used for holding single character values. Therefore, this data type cannot be used to hold strings or arrays of characters and does not take a size qualifier.  Corresponds to the primitive Java type <code>char</code> .  Fields of type SVR_CHAR cannot be null on the database.
SVR_DATE	Used for holding date values with a resolution of one day.  Corresponds to class <code>curam.util.type.Date</code> .  Fields of type SVR_DATE can be stored as null on the database.
SVR_DATETIME	Used for holding date and time values with a resolution of 1 second.  Corresponds to class <code>curam.util.type.Date</code> .  Fields of type SVR_DATETIME can be stored as null on the database.

Table 5. *Cúram Data Types (continued)*

Type	Description
SVR_DOUBLE	<p>Used for holding floating point numbers.</p> <p>Corresponds to the primitive Java type double.</p> <p>Fields of type SVR_DOUBLE cannot be null on the database.</p>
SVR_FLOAT	<p>Used for holding floating point numbers.</p> <p>Used for holding floating point numbers.</p> <p>Corresponds to the primitive Java type float.</p> <p>Fields of type SVR_FLOAT cannot be null on the database.</p>
SVR_INT8	<p>An 8-bit integer.</p> <p>Corresponds to the primitive Java type byte.</p> <p>Fields of type SVR_INT8 cannot be null on the database.</p>
SVR_INT16	<p>A 16-bit integer.</p> <p>Corresponds to the primitive Java type short.</p> <p>Fields of type SVR_INT16 cannot be null on the database.</p>
SVR_INT32	<p>A 32-bit integer.</p> <p>Corresponds to the primitive Java type int.</p> <p>Fields of type SVR_INT32 cannot be null on the database.</p>
SVR_INT64	<p>A 64-bit integer.</p> <p>Corresponds to the primitive Java type long.</p> <p>Fields of type SVR_INT64 might be null on the database.</p>
SVR_MONEY	<p>A fixed-point numeric value with two decimal places used for holding currency values.</p> <p>Corresponds to the primitive Java type <code>curam.util.type.Money</code>.</p> <p>Fields of type SVR_MONEY cannot be null on the database.</p>

Table 5. *Cúram Data Types (continued)*

Type	Description
SVR_STRING	<p>Used for holding string values.</p> <p>Corresponds to the Java class <code>java.lang.String</code>.</p> <p>A SVR_STRING might optionally have a length qualifier. A SVR_STRING without a length qualifier is a SVR_UNBOUNDED_STRING. Strings that are stored on the database must have a length qualifier to enable a maximum size to be specified for the database column.</p> <p>A SVR_STRING can be stored on the database as either CHAR, VARCHAR or CLOB depending on its size and the type of database.</p> <p>Fields of type SVR_STRING might be null on the database.</p>
SVR_UNBOUNDED_STRING	<p>Used for holding string values for which a maximum length does not need to be specified.</p> <p>Corresponds to the Java class <code>java.lang.String</code>.</p> <p>SVR_UNBOUNDED_STRING is the only Cúram data type that cannot be used by an attribute of an entity class. The SVR_UNBOUNDED_STRING data type does not allow the developer to specify its maximum size and cannot be used to define a database column. To define a string field on an entity, you must use SVR_STRING with a length qualifier.</p>

## Packages

The package structure in the UML meta-model does not affect any of the generated outputs. The hierarchy of the meta-model is effectively flattened during the build process.

The structure of the hierarchy is significant because options that are set at the package level apply to all classes and other packages within that package. Any option can be overridden in any of the subpackages by setting the option at that level to a different value.

## Options

### CODE\_PACKAGE

It is possible for two or more process or struct classes in the model to have the same name. Equally named classes are distinguished (on the server side only) by their CODE\_PACKAGE value, which might be specified for one of its containing packages.

The CODE\_PACKAGE option, when specified, effects struct, entity, facade, and process classes within that package and in the packages contained within that package. Applying the CODE\_PACKAGE option to a class moves that class into a package within the default package, curam, and includes any of the package's parent CODE\_PACKAGE options. The following example outlines how the CODE\_PACKAGE option works:

For example, the UML meta-model class `MyProcess` in the model causes the following Java classes to be created:

- `<ProjectPackage>.intf.MyProcess`
- `<ProjectPackage>.base.MyProcess`
- `<ProjectPackage>.fact.MyProcessFactory`

The developer must implement:

- `<ProjectPackage>.impl.MyProcess`

If the developer wants to create another class that is named `MyProcess`, they can do so if they create the class within a package for which a different `CODE_PACKAGE` option was specified. Creating a `MyProcess` class ensures that the corresponding Java classes can be stored in separate locations on disk.

The developer specifies the following option for the package that contains the `MyProcess` class (the option must be manually typed into the documentation for the package in the UML meta-model):

- `CODE_PACKAGE = custom`

The following classes and interfaces result:

- `<ProjectPackage>.custom.intf.MyProcess`
- `<ProjectPackage>.custom.base.MyProcess`
- `<ProjectPackage>.custom.fact.MyProcessFactory`

The developer must implement `<ProjectPackage>.custom.impl.MyProcess`.

#### **Rules for the `CODE_PACKAGE` Feature:**

- `CODE_PACKAGE` values must be valid Java identifiers.
- Setting the `CODE_PACKAGE` option for a package recursively affects subpackages and process, facade, entity, and struct classes within the package.
- If you specify a `CODE_PACKAGE` value in a package whose parent also specified a `CODE_PACKAGE` value, the value that you specified overrides the value specified by the parent.

For example:

- Package A contains package B
- Package A specifies `CODE_PACKAGE = cp1`
- Package B specifies `CODE_PACKAGE = cp2`

Then:

- The effective code package of classes in package A is `cp1`
- The effective code package of classes in package B is `cp2` (not `cp1.cp2`).
- A `CODE_PACKAGE` setting of `.` (dot) or `$` is interpreted as blank because a literal blank is ignored by the generator, and cannot be used to override a non-blank setting.
- Multiple level code packages can be specified by using similar syntax to Java packages whereby each level is delimited by a dot. For example, the following code package setting represents three levels of Java packages:  
`CODE_PACKAGE = cp1.cp2.cp3`
- The `CODE_PACKAGE` option allows multiple struct and process classes to have the same name, but only one instance of each facade class name can exist. Cúram clients currently cannot distinguish between multiple facade classes with the same name, regardless of their `CODE_PACKAGE` setting.

- Like process and struct classes, the behavior of the `CODE_PACKAGE` option with entity classes results in generated interface and struct classes that are produced in different packages. However, entity class names must still be unique throughout the application regardless of the `CODE_PACKAGE` option setting because all entities correspond to tables in the single underlying database.
- Generated list wrapper structs (triggered by the existence of readmulti operations) are produced in the same code package as the structs that they wrap. This code package might not be the same code package as the operation that caused their creation.

---

## Audit Mappings Classes

Audit fields contain extra information about the modification history of each record. They can be added to database tables for auditing purposes.

Audit fields are available on entity and struct classes and are updated only by certain entity operations.

Information in audit fields can be set by the developer and should include:

- Creation time
- Modification time
- Program ID
- User ID.

Audit fields consist of all the attributes of a special class in the input meta-model called `AuditMappings`. A field corresponding to each attribute of this class can then be automatically added to the database table, and also to all the standard details structs for the entity.

## Rules

The following rules apply to the `AuditMappings` class:

- The stereotype must be `audit_mappings`.
- The attributes of the class must be valid domain definitions.
- The class must be “flat”, that is, it cannot aggregate any other classes.

Audit mappings are made available to an application by adding a class named `AuditMappings` with a `audit_mappings` stereotype to the model. Individual entity classes can then enable audit mappings by setting the `Audit Fields` option.

If the meta-model contains an `AuditMappings` class, then provide a Java implementation class for it in the `impl` package.

**Note:** If this implementation class is not present, the server application cannot be compiled, and either:

- Delete the `AuditMappings` class from the model
- Explicitly disable audit mappings completely by specifying the generator switch `-noauditmappings` option.

The following rules apply to the `AuditMappings` implementation class:

- The class must contain the same fields as defined in the meta-model, that is, they must have the same name and data type.
- The fields must be public

- The class does not need to inherit from any other class
- Optionally, the class can contain the public void `set(final boolean isInsert, final boolean isModify)` method.

The public void `set(final boolean isInsert, final boolean isModify)` call-back method that is called whenever necessary (such as during inserts and modifies) by the data access layer. Do not use this call back method to populate the fields of the `AuditMappings` class. The two boolean parameters indicate whether the database operation is an insert or modify.

- Optionally, the class can contain a public void method that is named `set`, which takes no parameters. The data-access layer calls this method whenever fields need to be updated. Any public methods whose names start with `set` and take no parameters are called in arbitrary order. Support for using multiple setter methods will be discontinued.

If the details struct contains any of the audit mapping fields, then these fields are updated in the struct automatically during the operation and are included in the update or insert.

The corresponding fields of the audit mapping fields, which are not present in the details struct are updated on the database. That is, it is not necessary to include the audit mapping fields in the details struct for them to be updated on the database. However, the audit mapping fields are not included in table level auditing.

## Outputs

Turning on auditing for an entity has the following effect:

- Fields are automatically added to the entity and to the generated standard details struct for the entity.
- Infrastructure data-access code automatically makes calls to the `AuditMappings` class to populate its fields whenever audit fields are being updated on the database.

The following operation stereotypes cause audit information to be set:

- `modify`
- `nsmodify`
- `insert`
- `nsinsert`
- `nkmodify`
- `batchinsert`
- `batchmodify`

## Options

Two options are available for attributes of the `AuditMappings` class in the model:

- Exclude from insert
- Exclude from modify

If the Exclude from modify option is set for an audit mappings field, then the value of this field is not changed by a `modify`, an `nsmodify`, or an `nkmodify` operation. That is, the field is set when a record is inserted, and is never changed by subsequent updates. Similarly, if the Exclude from insert option is set, then the value of the field is not set by an `insert` or `nsinsert` operation but changed by any subsequent updates. The default value for each of these options is `false`.

**Note:** It is not possible to cause audit mapping fields to be excluded from operations of the ns stereotype. Handcrafted SQL in these operations can still be used to access audit mapping fields directly.

**Note:** If your audit mappings include a time stamp, then populate this field with the value returned by `TransactionInfo.getProgramTimeStamp()`. This value ensures that all audit mapping-enabled tables that are modified during the transaction have the same time stamp value even though they were not written to at the same time.

---

## Domain Definition Classes

Domains are data type definitions that resolve to a primitive data type or another domain.

In relational database terminology, a domain defines the range of values that are allowed for an attribute of an entity. Domain definitions work similarly in Cúram. Equivalent primitive types are supported across client, middleware, server, and database components of a Cúram application:

*Table 6. Domain primitive types at different levels of a Cúram application*

Cúram Architecture Layer	Datatypes
Server Remote Interface Layer	Java datatypes
Server Business Object Layer	Java datatypes
Server Data Access Layer	Java datatypes
Database	Database datatypes

By working with domains, rather than primitive types, developers are protected from having to worry about different representations of data in the various application layers. For this reason, entity and structure attributes must be defined in terms of a previously defined domain - it is not possible to use primitive data types directly.

Validations on each domain type are also allowed to be defined in the client application. A specific validation can then be executed for all attributes that are defined in terms of a domain type, before transactions are invoked on the server. This client-side pre-flight validation gives the user feedback on basic data type validation without having to call the server. This results in lower network traffic because of the reduced number of failed transactions.

## Options

The following are the options allowed for Domain Definitions:

### Code Table Name

This specifies the name of the code table which contains valid entries for this Domain Definition. If the domain definition represents a hierarchy of code tables, the name of the lowest code table in the hierarchy should be specified as the code table name.

For fields for which a code table has been specified, the client application will display a drop-down list of valid values for the field if it is editable, or the code table translation for the field if it is read-only. In the case of a code table hierarchy, if the code table field is editable, n-levels of drop-down lists are displayed, where n is the number of code tables in the hierarchy. Only the first level is populated and a selection must be made to populate the next level in the hierarchy. For a read-only field where the code table is a hierarchy, the translation for the lowest level code table only is displayed.

This option is only valid for Domain Definitions which have been defined in terms of one which has theCode Table Root option set to yes.

### **Code Table Root**

Specifies whether the current Domain Definition is the root of a hierarchy of code table Domain Definitions. If this is set to yes, then all Domain Definitions which use this one (i.e. they are defined in terms of it) must specify theCode Table Name option. If the developer forgets to specify theCode Table Name option, an error will be displayed by the generator.

Since this Domain Definition will be used to hold code table codes, its type should match that of a Cúram code table code, i.e. SVR\_STRING<10>.

For more information see “Code Table Name” on page 12 above.

### **Compress Embedded Spaces**

Implemented in the Cúram client application.

It specifies that any *extra* whitespace<sup>2</sup>(not all whitespace) embedded in the string, and that all leading and trailing whitespace is removed before being sent to the server.

*Extra* whitespace consists of a run of whitespace characters immediately *after* another whitespace character. This means that each run or sequence of whitespace characters is deleted except for the first whitespace character of the run. For example, a pair of words separated by three spaces will be converted to the pair of words separated by one space.

Note that in cases where the first whitespace character is not a space, the results may not be as expected. For example, a pair of words separated by carriage-return, line-feed, space, space will be converted to the pair of words separated by the carriage-return character.

Note also that if this feature is used on multiple line text fields it will remove indentation.

**Note:** Switching on this option also causes leading and trailing whitespace to be trimmed from the string, regardless of theRemove Leading Spaces andRemove Trailing Spaces option settings.

### **Convert to Uppercase**

Implemented in the Cúram client application.

---

2. A whitespace character consists of any character for which java.lang.Character.isWhitespace(char) returns true. Such characters include the space character, the tab character and the line-feed character.

It specifies that the contents of this string field be converted to uppercase before being sent to the server.

### **Custom Validation Function Name**

Domain Definition validations implemented in the client infrastructure include a custom validation type which corresponds to a developer-supplied function for performing validations on data entered by users via the client interface.

This option allows the developer to specify the name of this function which associates it with the application UML model. The value of the option should be simply the name of a function (just function, not class + function, since the class name is defaulted in the client code). It must also be a valid Java identifier.

**Note:** This feature has been deprecated, please see the “Custom Data Conversion and Sorting” chapter of the Web Client Reference Manual for information on the new domain plug-in system.

### **Default**

Implemented in the Cúram client application.

It specifies that this field will contain a default value after it is displayed.

### **Maximum Size**

Implemented in the Cúram client application and database DDL.

It specifies a maximum number of characters which can be entered to this field before it can be sent to the server and forms the field storage size on the database.

### **Maximum Value**

Implemented in the Cúram client application.

It specifies a maximum permitted numeric value which must be entered into this field before it can be sent to the server.

### **Minimum Size**

Implemented in the Cúram client application.

It specifies a minimum number of characters which must be entered to this field before it can be sent to the server.

### **Minimum Value**

Implemented in the Cúram client application.

It specifies a minimum permitted numeric value which must be entered into this field before it can be sent to the server.

### **Multibyte Expansion Factor**

Implemented in the Data Manager for DB2 and DB2 for z/OS only.

For string domains it specifies an expansion factor (float from 1.0 to 4.0) to be applied when multibyte character set (MBCS) data will be used with DB2 or DB2 for z/OS. It overrides the global build-time property (curam.db.multibyte.expansion.default.factor) and is only necessary in order to deviate from the global setting (e.g. a particular domain is causing a DB2 limit to be exceeded). A setting of 1.0 effectively turns off expansion for this domain. You might choose to set this option for domains where you know the contents will never contain localized data; e.g. they are constrained to programmatically-defined

Western characters and can't be input via a client. The same option set for a string entity attribute can override this domain setting (see "Multibyte Expansion Factor" on page 30 for more information). This option is ignored if the feature is turned off (`curam.db.multibyte.expansion` property set to `false` at build time, see the "Planning for MBCS Data" topic in the *Cúram Server Developer's Guide* for more information).

### **Pattern Match**

Implemented in the Cúram client application.

It specifies a regular expression that the string value must match before it can be sent to the server. The regular expression must match the whole string, not just a portion of it. The regular expression syntax is the standard Java regular expression syntax used in Java 1.5. Full details on the supported syntax for these regular expressions can be found in the JavaDoc documentation for the `java.util.regex.Pattern` class supplied with your Java SDK.

### **Remove Leading Spaces**

Implemented in the Cúram client application.

It specifies that any leading spaces be stripped off the string before it is sent to the server.

### **Remove Trailing Spaces**

Implemented in the Cúram client application.

Specifies that any trailing spaces be stripped off the string before it is sent to the server.

### **Storage Type**

It allows the developer to specify what type of string storage datatype to use for this domain definition on the database. This option is only relevant for string domain definitions for which a length has been specified.

## **Overriding a Domain Definition**

The SDEJ provides the facility to override existing domain definitions without modifying the original domain definition. This is desirable in situations where the original domain definition is provided by a third party and should not be modified locally.

Suggested uses:

- Change the maximum size of a string field.
- Change the `Storage Type` of a domain definition.

### **How to use Domain Definition Overrides**

A domain definition is overridden by creating a new domain definition with the same name prefixed by an asterisk. For example, the domain definition, `PERSON_NAME`, would be overridden by creating a domain definition named `*PERSON_NAME`. At build time the overridden version is used instead of the original version, complete with its own data type and options.

### **Considerations / Limitations**

- It is important to be aware that overriding a domain definition affects all usages of the original domain definition. It is the responsibility of the developer to ensure that pre-existing functionality is not broken by overriding domain

definitions. Specifically attempting to change the Type of the domain definition, the Code Table Name or the Code Table Root is discouraged.

## Usage Rules

- A domain definition may be overridden by only one override.
- A domain definition override cannot be overridden. For example, if PERSON\_NAME is overridden by \*PERSON\_NAME, it is *not* permitted to further override \*PERSON\_NAME with \*\*PERSON\_NAME.
- It is not possible to create overrides for domain definitions which do not exist. For example, if there is a domain definition override named \*PERSON\_NAME then the model must contain a domain named PERSON\_NAME.
- Domain definition overrides cannot be used as attributes of structs or entities, i.e. attributes cannot use domain definitions whose names begin with an asterisk.

---

## Entity Classes

An entity is a collection of fields and associated database operations. Entity classes are the fundamental building blocks of systems that are developed with Cúram. They correspond to database tables. The Cúram generator supports automatic code generation for entity classes.

Entity classes have a stereotype of entity. An entity class is essentially an object wrapper for a database table. The attributes of an entity are transformed to columns on the database table. Entities can have various data maintenance operations such as read, insert, modify, remove, readmulti (read multi reads multiple records from a table based on a partial key).

Standard operations such as read or insert operate on a single database table by default.

Entities can have attributes, operations, dependencies, inherits relations, and aggregations. Each of these constructs has a set of rules that are associated with it.

### Related reference:

“Data Access Layer” on page 3

## Rules

- Entities must have at least one attribute *unless* the entity is a subclass of another entity, in which case it must have no attributes.
- Entities are not allowed to aggregate other classes.

## Attributes

Entity attributes correspond to columns with the same name on their associated database table.

Attributes are not contained in the generated BOL or RIL. This is because Cúram interface objects are stateless and atomic. Instead these attributes are contained within generated standard key and details structs (see “Outputs” on page 28).

The stereotype of an entity attribute cannot be blank. It must be one of the following:

## Details

The attribute is included as a column on the database table and in the standard details struct for the entity. For more information, see “Standard Details Structs.”

## Key

The attribute is included as a column on the database table, it forms part of the primary key, it is included in both the standard details struct and the standard key struct for the entity. For more information, see “Standard Key Structs.”

## Operations

Entity operations can be divided into two categories as determined by their stereotype:

### Database Operations

These are operations whose stereotype is recognized by the generator. These operations are fully or partially generated by the generator and operate directly on the RDBMS table related to the entity. They include standard operations to *read*, *insert*, *update*, *delete*, together with their variants.

### Non-database Operations

These are operations whose stereotype is not recognized by the generator. The generator generates only prototypes and skeletons for these operations, no data-access operations are generated. The body of these functions must be implemented in the BOL by the developer.

The operations available for entity classes are listed in “Operation Stereotypes” on page 4.

## Outputs

Entity classes are transformed into classes with operations and no attributes. The attributes from the entity in the input meta-model are transformed into one or more structs.

### Standard Key Structs

Standard key structs are generated for entity classes and contain those attributes in the class whose stereotype is key. If no such attributes exist in the class then a standard key struct is not generated.

This struct will be used as a parameter for operations requiring a primary key. For example, reads and deletes.

Though standard key structs do not appear in the input meta-model they can be used as arguments to operations in the input meta-model. The name given to standard key structs is the name of the corresponding entity with the word Key appended. For example, the standard key struct for the class Employer would be called EmployerKey.

### Standard Details Structs

Standard details structs are generated for all entity classes and contain all the attributes of the class. This struct is used as a data parameter to insert, reads and updates. Structs containing arrays of standard details structs are returned from standard readmulti operations.

Though standard details structs do not appear in the input meta-model they can be used as arguments to operations in the input meta-model.

The name given to standard details structs is the name of the corresponding entity with the word `Dtls` appended. For example, the standard details struct for the class `Employer` would be called `EmployerDtls`.

### **Standard List Structs**

Standard list structs are generated for entity classes which contain one or more operations of stereotype `readmulti` or `nkreadmulti`. This struct contains a single attribute named `dtls` which is a sequence of the standard details struct for the entity.

The name for a standard list struct is the name of the standard details struct for the entity with the word `List` appended. For example, the standard details struct for the class `Employer` would be called `EmployerDtlsList`.

## **Options**

The options available for entity classes are described in the sections below.

### **Abstract**

Specifies that the class is abstract. Abstract classes are intended to be subclassed by other classes. For more information on abstract classes and subclassing, please see “Subclassing” on page 87.

### **Allow Optimistic Locking**

Only applicable for entities which are not subclasses. Optimistic locking is supported on certain database operations (see “Optimistic Locking” on page 36 in “Operation Options” on page 30). In order to use optimistic locking on an entity's operation, this option must first be switched on for the class.

For more information on optimistic locking, see “Concurrency Control - Optimistic Locking” on page 19.

### **Audit Fields**

Only applicable for entities which are not subclasses. Extra fields can be configured to store additional information on a database table for auditing purposes. These fields are covered in more detail in “Audit Mappings Classes” on page 10.

If this option is switched on, then the available pre-configured audit fields will be automatically added to this entity and its standard details struct.

### **Enable Validation**

The validation operation is an exit point which gets called automatically for the purpose of validating data. This exit point will get called before the data-access layer entity operations whose stereotype is `insert` or `modify`.

For more information on exit points, see “Exit Points” on page 21.

### **Last Updated Field**

Only applicable to entity classes that are not subclasses.

In order to use the last updated field feature for an entity class, this option must first be switched on. This results in the addition of an extra timestamp field to the specified entity. This field gets updated with the current date and time whenever the record is written - unless the write was performed by an `ns` operation. For more information on the last updated field feature, see “Last Updated Field” on page 23.

## No Generated SQL

Switches on theNo Generated SQL for all database operations of the entity class. Individual entity operations can override the value of this option.

For more details see “No Generated SQL” on page 35 in “Operation Options” on page 30.

## Replace Superclass

This is only relevant to entities which are subclasses.

If this option is set, then requests to create instances of the superclass will instead result in the creation of the subclass. This enables the developer to change functionality by replacing subclasses with other classes.

## Concurrency Control - Optimistic Locking

Using optimistic locking for concurrency control means that more than one user can access a record at a time, but only one of those users can commit changes to that record. Once one user has modified the record, another user cannot modify it without first re-reading the latest version of the record. Thus it is optimistic in the sense that one user does not expect another to attempt to modify the same record at the same time.

The record being edited is locked for update only while the changes are being committed. This has the advantage of minimizing the time for which a lock is in place.

The disadvantage of optimistic locking is that when a user begins to edit a record, they cannot be sure that the update will succeed. An update that relies on optimistic locking will fail if another user has updated a record while the first user is still editing it.

Optimistic locking is implemented by adding an extra field to the database table. The extra field contains the version number for the record and is automatically incremented each time the record is modified. The generated DAL code checks this version number while the record is being updated, and if the version number on the database table is not the same as the version number on the original record then the update operation is aborted and an exception is thrown.

Optimistic locking is permitted only on Entity classes.

The following operation stereotypes support optimistic locking:

- modify;
- nkmodify;
- nsmodify.

The following operation stereotypes are affected by optimistic locking:

- insert - The version number field is automatically included in the details parameter and is automatically initialized before being written to the database;
- nsinsert - If optimistic locking is enabled on an entity class, the version number field *must* be included in the details struct by the developer and will be automatically initialized before being written to the database.

Optimistic locking is only possible for operations which modify a single database record and whose details struct includes the generated Version Number field. This

means that for non-standard operations, it is up to the developer to ensure that the non-standard key parameter always identifies a single unique record and that the Version Number field is included in the details struct. For `nkmodify` operations, optimistic locking is only possible if the database table contains exactly one record. This field must be called `versionNo` and its type should be `VERSION_NO`. The developer must ensure that the model contains a numeric domain definition named `VERSION_NO`.

In order to support optimistic locking on an operation you must do two things:

- Switch on the `Allow Optimistic Locking` option on the entity.  
This will cause the Version Number field to be automatically added to the entity.
- Switch on the `Optimistic Locking` option on the operation.  
This will cause the generator to generate code in the DAL for the operation which will check and update record version numbers accordingly.

## Table Level Auditing

Auditing is supported on all stereotyped entity operations except `ns`, `nsmulti`, `batchinsert` and `batchmodify`.

The information captured by table level auditing is stored in the database table `AuditTrail`.

Table level auditing is enabled by switching on the `Database table-level auditing` option for an operation. This causes the generated data-access code to record audit information for an operation.

The type of audit information recorded depends on whether optimistic locking is switched on or off for the operation. If optimistic locking is switched on, then the audit information includes the information of the new and old versions of the record, otherwise it only includes information about the SQL operation invoked.

### Information Captured by Table-level Auditing

The following information is captured:

- *Date and time* - The date and time of the transaction.
- *User ID* - The ID of the user who invoked the transaction.
- *Table name* - The name of the database table which was modified.
- *Program name* - The FID of the function which invoked the transaction.
- *Transaction type* - Indicates whether the transaction was *online* / *batch* / *deferred* / etc.
- *Key info* - The key which was provided to this operation. Note that this may identify one or many records.
- *Details of changed data* - These details are logged in an XML format. The exact format of this XML can be seen in the JavaDoc details for the class `curam.util.audit.AuditLogInterface` in the `doc/api` directory of the SDEJ. They include the names of the all the fields referenced by the details struct, the field types, the new version of the field data and, if optimistic locking is enabled, the old version of the field data.

If optimistic locking is switched on, then the operation is guaranteed to have only affected a single record. Therefore, the audit information includes information about the record before and after the operation. The old version of the record is re-read, the old value of each field is compared to the new value, and any field which has changed is included in the audit information, i.e. unchanged fields are filtered out.

If optimistic locking is switched off then for performance reasons the record is not re-read during the update, so the audit information will contain only the new versions of all the fields involved in the update, not a *before-after* comparison of the record. Also, any non-optimistic updates apart from the 'modify' stereotype can potentially affect more than one record, in which case it is not possible to record a before-after comparison of the update. All the detail fields will be included regardless of whether the new value is different to the old value.

This data can be compressed when using the default auditing handler by specifying the `curam.audit.audittrail.datacompressionthreshold` property described in the *Cúram Configuration Parameters*.

- *Operation type* - Indicates whether the operation was one of: create, read, update or delete.

### **Storage of Audit Information**

By default, the audit information captured is written to the AuditTrail database table. The developer may also supply their own auditing handler by specifying a class which implements the `curam.util.audit.AuditLogInterface` interface. For more information, see the section on Customization Settings in the *Cúram Server Developer's Guide*.

## **Exit Points**

An exit point is a callback function written by the developer and executed at a predefined strategic point by the server.

Four types of exit point are supported:

### **Pre Data Access**

This function is called before the DAL function (but after *Validate* functions).

The function is named after the method to which it belongs, prefixed with *pre*, e.g. *preread*.

### **Post Data Access**

This function is called after the DAL function.

The function is named after the method to which it belongs, prefixed with *post*, e.g. *postread*.

### **Validation**

This function is called before standard insert and standard update operations, and also before *Pre-data Access* functions. It provides a common place to put validation code.

The function is named *autovalidate*. Note that this exit point is enabled per entity rather than per operation.

The validation exit point always has exactly one parameter which is the standard details struct for the entity, and is declared to throw the same exceptions as stereotyped operations of the entity.

Since it is only insert and modify which are guaranteed to pass in the standard details struct, it is only these operation stereotypes which can utilize the validation exit point. Other operation stereotypes do not utilize this exit point, even if they have the standard details struct as one of their parameters.

## On-fail

This function is called if an error occurs in the data access function.

The function is named after the method to which it belongs, prefixed with `onFail`, e.g. `onFailread`.

**Note:** For non-void operations the return class is included in the arguments to this method and will always be *null*.

## Exit Point Parameters

With the exception of *Validate* exit points whose parameters are described above, the parameters to an exit point method consist of the following:

- the parameters to the method to which the exit point belongs. (In fact if any extra parameters have been specified for a database operation in the model, this is the only place where the developer can access them.)
- the return type of the method to which the exit point belongs - if a return type is present.

**Limitation:** The return type parameter will not be included into the parameters of exit point methods for `nsread` and `ns` operations.

The following approach can be used to generate the return type parameter into the parameters of exit point methods for `nsread` and `ns` operations:

- Add an un-stereotyped method to the entity class giving it the same signature as the `nsread` or `ns` operation.
- Set the *Post Data Access* option on your `nsread` or `ns` operation to `False`.
- The implementation of your un-stereotyped operation will then call the `nsread` or `ns` operation, and will have access to its return value as required.
- for on-fail exit points, an exception class. This is the exception which was thrown from the data access layer. The exit point may handle the error or pass it on by throwing it.

## What should exit points be used for

Exit points are intended to be used for validation or for completing a business process. For example, after modifying an invoice detail line record, the last modified date should be updated on the invoice header record.

## What should exit points *not* be used for

It is recommended that exit points should not be used as a means of populating incomplete fields in incoming parameters. This situation should be handled by wrapping the database function in a non-database function which would take a copy of the incomplete record, fill in the missing fields and invoke the database operation.

When adding an exit point to an entity operation, ensure that this will not have any side effects for other users of the operation.

## Entity Inheritance

Input meta-model Entity classes are allowed to subclass other entity classes. Typically, entity classes will be subclassed in order to add functionality (such as additional stereotyped operation) which are required for special processing of the associated database table, but which do not belong in the parent class.

For more information please see “Subclassing” on page 87.

## Rules when Using Entity Inheritance

- Entity classes are only allowed to inherit from other entity classes.
- Subclasses of entities can add any number of additional database and user-defined operations.
- Subclasses of entities cannot add attributes. This is because the underlying relational database table must not be affected by the inheritance.
- Entity subclasses do not have standard key and details generated for them, they use the standard key and details structs from the base class.

## Last Updated Field

This feature is similar to the Audit Mappings feature. It is a field which can be added to database tables to contain extra information about the modification time of each record for reporting purposes.

The feature is only available for entity classes and it is updated only by certain entity operations

Switching on the last updated field functionality for an entity has the following effects:

- A field called *lastWritten* of type SVR\_DATETIME is automatically added to the entity.
- The Cúram infrastructure automatically populates this field with the current time whenever the record is written to the database - unless the write was performed by an ns operation.

The following steps must be taken to avail of this feature:

- To turn on the feature for an individual entity class the `Last_Updated_Field` property in the Rational Software Architect Cúram Properties tab must be set to '1 - yes' using the supplied drop-down.
- To turn on the feature for all of the entities for a particular application, the following text must be appended to the `extra.generator.options` property in the `Bootstrap.properties` file as follows:

```
extra.generator.options=-defaultoption class_lastupdatedfield=yes
```

Figure 1. `extra.generator.options` property in `Bootstrap.properties`

```
extra.generator.options=-defaultoption class_lastupdatedfield=yes
```

- A new domain definition must be specified in the model as follows:
  - Domain Definition Name: `LAST_UPDATED`
  - Domain Definition Type: `SVR_DATETIME`

Invoking operations with the following stereotypes cause the *lastWritten* field to be set:

- `insert;`
- `nsinsert;`
- `modify;`
- `nkmodify;`
- `nsmodify.`
- `batchinsert.`
- `batchmodify.`

**Note:** Unlike the version number field utilized for the optimistic locking feature, there is no requirement to add the last written field to structures involved in non-standard insert and modify operations. If the last updated field feature has been enabled for an entity, this field is always updated for the operation stereotypes listed above by the infrastructure data access code regardless of whether the field is present in the structure being used.

---

## Extension Classes

Extension classes allow you to specify options for a target class without modifying the meta-model definition of the target class.

Each extension class should be linked to one target class. At build time the contents of an extension class are effectively super-imposed on its target class.

### How to use Extension Classes

To extend an existing class, create a new class of stereotype Extension; see *Working with the Cúram Model in Rational Software Architect* for more information on using and modeling with Rational Software Architect.

Options can be added to the extension class in the same way as for other classes. When any of these are added to an extension class they have the effect of adding (if not already existing) or modifying (if already existing) the same named option on the target class.

When creating an Extension class in Rational Software Architect, keep in mind that since it can apply to different class types. You must make sure the settings for the extension class are compatible with the class you are extending.

The extension class is linked to its target class by adding an relationship of stereotype extension between the two classes. The new class should be created within a custom sub-package of your model.

### When to use Extension Classes

Extension Classes should only be used for the following purpose:

- To switch on theLast Updated Field option on an entity.
- To switch on theAudit Fields option on an entity or struct.

**Note:** The Rational Software Architect user interface allows you to specify additional information such as attributes and options for extension classes but only the above two options should be included in the class. Other changes are not compliant.

### Considerations / Limitations

- Storage of relationships. When a relationship is created between two classes in Rational Software Architect, it should be noted that the relationship is not stored within either of the actual classes but rather as a free standing object in a package. (Usually the relationship is stored in the package containing the diagram on which it was drawn but this is not guaranteed.) The developer should ensure that the relationship is stored in a location where it will not be lost or overwritten during an upgrade. Inheritance relationships are always stored within the subclass so there is no risk of inadvertently losing these.

## Usage Rules

- An extension class may only be applied to one target class.
- A class may be extended by multiple extension classes.
- Extensions may be applied to classes of stereotype entity, struct.

---

## Facade Classes

Facade classes encapsulate a business process that is visible to the client. It is a collection of operations. Facades are the business object layer of the application.

Facade classes do not have data maintenance operations, or any relationship with database tables. They manipulate other entity and process classes in order to implement a business process. Facade classes have a stereotype of facade.

## Rules

- Facade classes must have a stereotype of facade.
- Facade classes cannot have aggregations to any other classes.
- Facade classes can only inherit from other facade classes - not from entity or process classes.
- Facade classes cannot have attributes.
- Facade classes cannot have the same name. See “Options” on page 8 for more details.

## Operations

Within facade classes there are four operations supported:

### **default**

The default stereotype offers a standard or plain operation.

### **batch**

For operations of stereotype batch, the Cúram generator will produce the necessary source code wrappers to build a batch wrapper program which will enable this operation to be run by the Batch Launcher (see the *Cúram Batch Processing Guide* for more details).

The rules when defining batch operations are:

- Batch operations cannot have more than one parameter.
- Parameters to batch operations must be structs.
- A facade class cannot have more than one batch operation.

### **wmdpactivity**

A method of a facade class can be designated as a deferred processing method by setting its stereotype to wmdpactivity.

For more information please consult the *Cúram Server Developer's Guide*.

### **qconnector**

For operations of stereotype qconnector, the generator will produce the necessary source code to connect to a JMS provider (e.g. IBM MQSeries®). For more information please see “Cúram JMS Queue Connectors” on page 76.

## Options

### Abstract

Specifies that the class is abstract. Abstract classes are intended to be subclassed by other classes. For more information on abstract classes and subclassing, please see “Subclassing” on page 87.

### Generate Facade Bean

This causes a stateless session bean to be generated for this class. This bean class can be used to allow your server to be accessed by other systems or by message driven beans.

### Replace Superclass

This is only relevant if this facade class has been subclassed from another class. For more information on subclassing, please see “Subclassing” on page 87.

---

## Process Classes

Process classes are a collection of operations that encapsulate a business process. Process classes do not have data maintenance operations, or any relationship with database tables. They manipulate other entity and process classes to implement a business process.

For example, a banking system has an account transfer process which debits money from one account and credits another. Internally, the process uses an account entity to debit one account, and credit another account. The process class itself does not do any database manipulation. It packages a sequence of entity operations to carry out the modeled business process. Process classes have a stereotype of process.

### Related reference:

“Business Object Layer” on page 2

## Business Process Objects

*Business Process Objects (BPO s)* are the classes which reside in the *Business Object Layer (BOL)* of a Cúram server application, i.e. the architectural layer between the *Remote Interface Layer (RIL)* and the *Data Access Layer (DAL)*. All business logic is implemented in this layer and as a result, BPOs constitute the large majority of handcrafted coding required to create a server application.

BPOs do not *directly* communicate with the RDBMS (implemented - largely automatically - in the DAL), nor the middleware (implemented - largely automatically - in the BOL): their job is specifically to implement business logic.

## Rules

- Process classes must have a stereotype of process.
- Process classes cannot have aggregations to any other classes.
- Process classes can only inherit from other process classes - not from entity classes.
- Process classes cannot have attributes.
- Two or more process classes can have the same name provided that different `CODE_PACKAGE` values have been specified for each. See “Options” on page 8 for more details.

## Operations

Within process classes there are four operations supported:

### **default**

The default stereotype offers a standard or plain operation.

### **batch**

For operations of stereotype batch, the Cúram generator will produce the necessary source code wrappers to build a batch wrapper program which will enable this operation to be run by the Batch Launcher (see the *Cúram Batch Processing Guide* for more details).

The rules when defining batch operations are:

- Batch operations cannot have more than one parameter.
- Parameters to batch operations must be structs.
- A process class cannot have more than one batch operation.

### **wmdpactivity**

A method of a process class can be designated as a deferred processing method by setting its stereotype to wmdpactivity.

For more information please consult the *Cúram Server Developer's Guide*.

### **qconnector**

For operations of stereotype qconnector, the generator will produce the necessary source code to connect to a JMS provider (e.g. MQSeries). For more information please see "Cúram JMS Queue Connectors" on page 76.

## Options

### **Abstract**

Specifies that the class is abstract. Abstract classes are intended to be subclassed by other classes. For more information on abstract classes and subclassing, please see "Subclassing" on page 87.

### **Generate FIDs**

Specifies that the class should have a Function Identifier generated for it.

### **Replace Superclass**

This is only relevant if this process class has been subclassed from another class. For more information on subclassing, please see "Subclassing" on page 87.

---

## Struct Classes

Struct classes are Java classes with public attributes and no modeled methods. A struct is the Java equivalent of a C++ struct. Structs allow for the grouping of domain definitions and other struct classes to form programmatic record definitions.

Struct classes are used as arguments to operations of entity and process classes. Structs are used to package arguments to avoid long argument lists. Struct classes can also aggregate each other. These aggregations become struct members.

For example, a bank account entity has parameters to a read operation that consists of a key struct and a details struct. The key struct contains a single field

for the account number. The details struct contains several fields such as Name, Balance, and so on. Struct classes have a stereotype of struct.

## Rules

- Struct classes must have one or more attribute or aggregation, i.e. a struct cannot be empty.
- Struct classes are not allowed to have operations.
- Struct attribute types must be defined in terms of valid domain definitions.
- Struct classes may aggregate entity classes or other struct classes.
- Struct classes are not allowed to be involved in inheritance relationships.
- Struct classes used as key or details parameters to non-standard database operations must not aggregate other structs, i.e. they must be “flat”.
- Two or more struct classes can have the same name provided that differentCODE\_PACKAGE values have been specified for each, i.e. similarly named struct classes must be distinguishable by havingCODE\_PACKAGE settings. See “Options” on page 8 for more details.
- In most cases you will have to define a struct in order to use it as a parameter to an operation. The exception to this rule is standard key and details structs. These are generated automatically by the Cúram generator and are available for use by the developer.

## Outputs

Input meta-model struct classes map directly onto generated Java classes in the <ProjectPackage>.<CodePackage>.struct package. The Java struct class contains public fields corresponding to each attribute defined in the model.

Each field is initialized to its default value - zero for numerics, empty string for Strings, etc. - so the developer does not have to worry about null values.

Each field is accompanied by comments describing the Domain Definition hierarchy for the datatype.

The class also contains generated code enabling the struct to be cloned and assigned to other structs.

Struct classes have no counterpart in generated DDL.

## Options

### Audit Fields

If this option is switched on then the available pre-configured audit fields will be automatically added to this struct.

This option should be enabled if the struct class is being used as a write operation of an entity which also hasAudit Fields switched on.

For more information, see “Audit Mappings Classes” on page 10.

---

## Attributes

Attributes represent fields of the underlying Java class. The class stereotype determines the attribute stereotypes that are valid for the class.

The combination of class and attribute stereotypes determines how the Cúram generator processes the UML meta-model.

## Attribute Rules

The following table shows the mapping of class stereotypes to attribute stereotypes.

*Table 7. Mapping of Class and Attribute Stereotypes*

Class Stereotype	Valid Attribute Stereotypes
audit_mappings	audit_mappings
domain_definition	N/A
entity	details, key
facade	N/A
listrdo	dataitem
loader	N/A
process	N/A
rdo	dataitem
struct	default
webservice	N/A
wsinbound	N/A

- Attribute names must be unique within a class.
- Attributes must be defined in terms of domain definitions.
- Since attributes ultimately appear in generated Java code, their names must be valid Java identifiers.
- The order of attributes in the primary key of an entity is determined by the order in which the attributes appear in the entity class. Since their order in the entity is not critical, you can change this order to obtain the primary key configuration you desire.

## Attribute Options

### Allow NULLs

This option is available only for details stereotyped attribute on an entity class.

It determines whether NULL values are permitted on the corresponding database field. Setting this option to no causes a Not Null qualifier to be included with this field in the generated DDL script.

The default value for this option is dependent on the underlying data type of the field. The default value of this option for the attributes for fields of type SVR\_BOOLEAN, SVR\_CHAR, SVR\_FLOAT, SVR\_DOUBLE, SVR\_MONEY, SVR\_INT8, SVR\_INT16, SVR\_INT32 is no.

The default value of this option for the attributes of type SVR\_BLOB, SVR\_DATE, SVR\_DATETIME, SVR\_STRING, SVR\_INT64 is yes.

This topic is dealt with in “Null Considerations” on page 59.

## Multibyte Expansion Factor

This is an override for the domain-level Multibyte Expansion Factor, described in “Multibyte Expansion Factor” on page 14, and is applicable to string entity attributes only.

It specifies an expansion factor (float from 1.0 to 4.0) to be applied when multi-byte character set (MBCS) data will be used with DB2 or DB2 for z/OS. It operates in conjunction with its equivalent domain option and the global build-time properties `curam.db.multibyte.expansion.default.factor` and `curam.db.multibyte.expansion`, which are described in the *Cúram Server Developer's Guide*. This option is only necessary for DB2 MBCS data in order to deviate from the global or domain settings. For example, you might choose to set this option to 1.0 (which effectively turns off expansion) for a string attribute where you know the contents will never contain localized data (e.g. they are constrained to programmatically defined Western characters and can't be input via a client). This option is ignored if the feature is turned off via the `curam.db.multibyte.expansion` property (see the "Planning for MBCS Data" topic in the *Cúram Server Developer's Guide* for more information).

---

## Operations

Operations represent the functions of modeled classes that can be handcrafted or generated by the Cúram generator; depending on their type.

- Operations must belong to either entity, process, facade, web service, or wsinbound classes.
- Operations can be fully handcrafted or can use the features that are offered by standard operations.
- Operations cannot be individually hidden or exposed to clients, only whole classes can be hidden or exposed.

## Rules

These are the rules regarding the requirements for using structs (versus domain values) as parameters and return values for operations:

- Parameters for batch operations must be structs.
- Parameters and return types for all database operations must be structs.
- Parameters and return types for queue operations must be structs.
- Parameters and return types for web service connector operations must be structs.
- Parameters and return types for client-visible operations must be structs. (Domain parameters and return types are not supported by the HTML client.)
- Parameters and return types for other operation stereotypes including web service client operations or other classes may be domain definitions.

## Operation Options

### Audit BI (Business Interface) Calls to this Operation

This is only relevant to client-visible operations.

This option specifies whether Business-Interface-level auditing should be performed for this operation. For Business-Interface-level auditing records the following information about the operation call is recorded:

- the operation name (Function Identifier);

- the username of the caller;
- the date and time;
- the transaction type (online/batch/deferred/etc.).

This option can be overridden at application startup time using application properties, this functionality is described with an example in the Database Table-level Auditing option description contained in this section.

### **Auto ID Field**

This is only relevant for certain insert operations of entity classes.

Specify which field is to be used as the Auto ID field. The Auto ID field is automatically populated with a generated unique ID during the insert to ensure that the record can be uniquely identified.

### **Auto ID Key**

This option is used only in conjunction with the `Auto ID Field` option.

It allows you to specify the key set from which a unique ID should be generated.

### **Business Date**

This option is only relevant to operations of a process class.

It allows you to specify that one field of the operation parameters be treated as the Business Date Field for the operation. This means that the value of this parameter to the operation becomes the Business Date for the duration of the transaction. The Business Date is the `Date` or `DateTime` which gets returned by the following methods:

- `curam.util.transaction.TransactionInfo.getBusinessDateTime()`
- `curam.util.transaction.TransactionInfo.getBusinessDate()`
- `curam.util.type.Date.getCurrentDate()`
- `curam.util.type.DateTime.getCurrentDateTime()`

The main purpose of this feature is to give greater flexibility when running batch programs for which processing dates are significant. Consider the example where a report generating program is run at the end of each day to count all payments issued that day. The payment records are obtained by reading all records whose issue date equals `curam.util.transaction.TransactionInfo.getBusinessDate()`. This program will process a different set of records depending on the day on which it is run.

Now consider what would happen if you needed to re-generate the report from 10 days ago.

Without the `Business Date` feature you would have to do the following:

- Submit a batch request for your batch program.
- Change the system date on the machine where the batch program will be run. Note that you will have to ensure that this doesn't affect anyone else, so nobody else can use the machine while the system date is being changed.
- Ensure that your batch request is the only one in the queue.
- Run the batch launcher to cause your batch program to be run.
- Revert the system date on the machine.
- Make the machine available for general usage again.

However if your batch program parameters include a Business Date field you need only do the following:

- Submit a batch request for your program, ensuring that the batch job parameter which has been specified as the Business Date is set to the date 10 days ago.
- Run the batch launcher.

**Syntax for Business Date option::**

The Business Date option should be specified in one of the following formats:

- *fieldName*
- *paramName.fieldName*

where

- *paramName* is the name of a parameter. This is optional and, if not specified, the first operation parameter is assumed.
- *fieldName* is the name of a field in the parameter struct.

**Business Date example 1**

- Struct ReportArguments contains a Date field named *effectiveDate*.
- The following is a batch operation: `doReportGeneration (ReportArguments arg1)`.
- To use *effectiveDate* as the Business Date for the operation, you can set the Business Date option to either `arg1.effectiveDate` or because it is the first (and only) parameter: `effectiveDate`.

**Business Date example 2**

- Struct GeneratePaymentsParameters contains a Date field named *paymentDate*.
- The following is a batch operation: `generatePayments (SomeStruct argA, GeneratePaymentsParameters argB)`.
- To use *paymentDate* as the Business Date for the operation, you would set the Business Date option to `argB.paymentDate`.

**Rules for Business Date option::**

- This option is only relevant to operations which correspond to individual server transactions. Such operations are the operations of facade classes and batch operations. Note that it is *not* applicable to workflow activity or deferred processing operations.
- The field which is specified as the Business Date Field must be of type SVR\_DATE or SVR\_DATETIME.
- The Business Date Field only takes effect when the operation is invoked by a remote client (either the HTTP client or a web services client) or by the Batch Launcher. It does not take effect for operations which are invoked directly from Java code. This is because the latter does not result in a new server transaction being started.
- If the Business Date Field is set to null, `curam.util.type.Date.kZeroDate` or `curam.util.type.DateTime.kZeroDateTime` for a method invocation, it is ignored and the Business Date does not get overridden for that transaction. In this case the Business Date for the transaction will be either the current system date, or the overridden value specified in application properties - see the Date and DateTime JavaDoc documentation for more details.

**BytesMessage encoding character set**

This is only relevant for qconnector operations of process classes. See “Options on qconnector Operations” on page 77 for more information.

## Database Table-level Auditing

This is only relevant to database operations of entity classes.

This option specifies whether table-level auditing should be performed for this operation. Table-level auditing records detail information about the changes made to actual data on the database table.

The behavior of auditing depends on whether Optimistic Locking is switched on or off for the operation. For more information about Auditing, see “Table Level Auditing” on page 20.

This option can be overridden at application startup time using application properties, this functionality is available to Audit BI Calls and this option and what follows is an example of how it should be used.

**Changing operation auditing options without rebuilding.:** Changes to operation options Audit BI and Database Table-level auditing in the model require a rebuild and redeploy to take effect. It is possible to override these properties in application properties whereby the changes take effect when the application is restarted.

These two options can be targeted at individual operations by specifying application properties whose format is as follows:

```
curam.audit.audittrail.<ProjectName>.<ClassName>.<OperationName>  
curam.audit.opaudittrail.<ProjectName>.<ClassName>.<OperationName>
```

or, if the class is in a code package:

```
curam.audit.audittrail.<ProjectName>.<CodePackage>.  
<ClassName>.<OperationName>  
curam.audit.opaudittrail.<ProjectName>.<CodePackage>.  
<ClassName>.<OperationName>
```

Properties whose names begin with `curam.audit.audittrail` apply to the Database Table-Level Auditing option and cause data to be captured to table `AuditTrail`.

Properties whose names begin with `curam.audit.opaudittrail` apply to the Audit BI calls option and cause data to be captured to table `OpAuditTrail`.

Example (1): To switch on table level auditing for operation `modify` of entity `CaseHeader` which is in code package `core` of the `Cúram` application, set the property `curam.audit.audittrail.curam.core.CaseHeader.modify` to `true`.

Example (2):. To switch off operation auditing for operation `modifyAddress` of process class `Participant` which is in code package `core.facade` of the `Cúram` application, set the property

```
curam.audit.opaudittrail.curam.core.facade.Participant.modifyAddress
```

to `false`.

In Summary,

- changing the value of an auditing option requires an application restart to take effect
- The `curam.audit.opaudittrail.*` properties only affect client-visible operations.
- The `curam.audit.audittrail.*` properties only affect stereotyped entity operations - excluding stereotypes `ns` and `nsmulti`.

## Field Level Security

Field Level Security can be applied for the fields returned by client-visible operations (i.e. operations of the Facade class). It is only relevant to operations of a client-visible operation (defined in a Facade class).

In Rational Software Architect the **Secure Fields** properties tab of the Facade class operation allows you to apply security to any field returned by an operation by specifying a security identifier (SID) for that field.

To establish secure returned fields for an an operation use the **Secure Fields** button from the properties tab for the operation. Clicking the **SID Name** cell for the returned **Field Name** allows you to enter the security identifier (SID). The maximum length of a security identifier is 100 characters.

The client infrastructure will then ensure that fields for which a SID has been specified can only be viewed by users to whom that SID has been granted. Fields for which no SID has been specified will be visible to all users.

All the information about Field Level Security - which SID is assigned to a field - is written by the generator to an XML file and is loaded into database table FieldLevelSecurity by the Data Manager. The Data Manager configuration file datamanager\_config.xml must be changed to reference the generated file <ProjectName>\_FieldsReturned.xml. This can be done by adding an entry to the **initial** target as shown in "Field Level Security" below.

```
<target name "initial"
<entry
  name="build/svr/gen/ddl/<ProjectName>_Fids.xml"
  type="xml" base="basedir" />
<entry
  name="build/svr/gen/ddl/<ProjectName>_FieldsReturned.xml"
  type="xml" base="basedir" />
</target>
```

Figure 2. Sample datamanager\_config.xml for adding field level security information to the database

Once the field names and SIDs have been added to the FieldLevelSecurity table, the SIDs should be loaded into the SecurityIdentifier to enable them to be assigned to groups. This can be done using the database command shown in "Field Level Security" below.

```
INSERT INTO SecurityIdentifier(sidName, sidType, versionNo)
SELECT DISTINCT sidName, 'FIELD', 1 from FieldLevelSecurity
WHERE sidName IS NOT NULL;
```

Figure 3. Inserting field level security SIDs into the infrastructure SecurityIdentifier table

These SIDs can then be assigned to user groups using the Security Administration console.

## JNDI name of the QueueConnectionFactory class

This is only relevant for qconnector operations of process classes. See "Options on qconnector Operations" on page 77 for more information.

## JNDI name of the transmission queue

This is only relevant for qconnector operations of process classes. See "Options on qconnector Operations" on page 77 for more information.

## JNDI name of the reply queue

This is only relevant for qconnector operations of process classes. See “Options on qconnector Operations” on page 77 for more information.

## Message type

This is only relevant for qconnector operations of process classes. See “Options on qconnector Operations” on page 77 for more information.

## No Generated SQL

This is only relevant to database operations of entity class.

Switches off generation of data access code, allowing developers to provide their own implementation.

For example, if No Generated SQL was set to yes for a standard read operation named `myRead` the generator will produce a declaration of an abstract method named `myRead_da` with the same signature as the formerly generated `myRead` method. The developer must provide the implementation of method `myRead_da` as is shown in the following listing:

```
public MyEntityDtIs myRead_da(  
    final MyEntityKey key, final boolean forUpdate)  
    throws ApplicationException, InformationalException {  
  
    final MyEntityDtIs result = new MyEntityDtIs();  
    result.idNumber = "1234";  
    return result;  
}
```

Figure 4. Handcrafted data access implementation for a standard read

For `readmulti` operations - i.e. operations of stereotype `readmulti`, `nsreadmulti`, `nkreadmulti` or `nsmulti` - the handcrafted implementation must follow a different pattern. The method is declared as returning a list struct *but this return value is ignored*. `Readmulti` operations in Curam are implemented using the visitor design pattern whereby a subclass of `curam.util.dataaccess.ReadmultiOperation` is passed into the data access operation which then invokes its `operation(Object)` for each record found. Usually this operation will add the record to a collection which gets returned to the caller. This is described in greater detail in the *Curam Server Developers Guide*.

The key point is that for `readmulti` operations, data is returned to the caller by adding it to the `ReadmultiOperation` class by calling its `operation(Object)` method, and *not* by simply returning it from the method. This is shown in the following example:

```

/*
 * This implementation returns two hard coded dummy records.
 */
public MyEntityDtIsList readmulti_da(
    final SomeKey k, final ReadmultiOperation op,
    final boolean requireInformational)
    throws ApplicationException, InformationalException {

    // Create and add one record for return to the caller.
    final MyEntityDtIs oneDtIs = new MyEntityDtIs();
    oneDtIs.idNumber = "2222";
    op.operation(oneDtIs);

    // Create and add another record for return to the caller.
    final MyEntityDtIs twoDtIs = new MyEntityDtIs();
    twoDtIs.idNumber = "3333";
    op.operation(twoDtIs);

    // our return value is ignored so just return null.
    return null;
}

```

Figure 5. Handcrafted data access implementation for a readmulti

## On Fail Operation

This is only relevant for database operations of entity classes.

This option switches on the on-fail exit point.

If any error occurs in the Data Access Layer (DAL), this function is invoked with a copy of the parameters given to the DAL and a copy of the DAL exception corresponding to the error.

The type of exception depends on the type of error which occurred. The error can either be handled in this exit point or the exception can be thrown from here to allow the error to be handled elsewhere.

For more information on Exit Points, see “Exit Points” on page 21.

## Optimistic Locking

This is only relevant for certain update operations of entity classes.

This option switches on optimistic locking for this operation.

Note that this option is only allowable if theAllow Optimistic Locking option has been set for the entity class.

For more information on optimistic locking, see “Concurrency Control - Optimistic Locking” on page 19.

## Order By

This option applies only to entity operations of stereotype readmulti, nsreadmulti and nkreadmulti.

This option allows you to specify the fields by which a sequence of records are sorted as they are read from the database. Any or all of the fields of an entity are valid arguments for this option. Records are always sorted in ascending order.

If this option is not specified, records will be returned in arbitrary order.

### **Post Data Access Operation**

This is only relevant for database operations of entity classes.

This option determines whether a standard database operation has a post-exit point.

For more information on exit points, see “Exit Points” on page 21.

### **Pre Data Access Operation**

This is only relevant for database operations of entity classes and determines whether a database operation has a pre-exit point.

For more information on exit points, see “Exit Points” on page 21.

### **Readmulti\_Max**

This option applies only to entity operations of stereotype readmulti, nsreadmulti, nkreadmulti and nsmulti.

This allows you to specify the maximum number of records returned by a readmulti operation. If there are more records available than the Readmulti\_Max, then handling is based on the setting of the Readmulti\_Informational option (“Readmulti\_Informational”). Unless the Readmulti\_Informational setting is on in the model for the operation there is no Readmulti\_Max enforcement.

If this option is not specified then the generator system default will be used.

Specifying a value of 0 for this option is interpreted as infinity and no limit will be applied to the number of records returned.

### **Readmulti\_Informational**

This option applies only to entity operations of stereotype readmulti, nsreadmulti, nkreadmulti and nsmulti.

This allows you to determine the handling of the system when the specified Readmulti\_Max is reached (“Readmulti\_Max”). The default behavior is that a Readmulti\_Max message gets logged and all entries are returned to the user. If this option is specified, then an InformationalMessage can be added to the current transactions InformationalManager, for handling in the Application's Facade layer. In this case only the specified Readmulti\_Max number of entries will be returned to the user.

If this option is not specified then the generator system default false will be used.

### **Response message timeout (seconds)**

This is only relevant for qconnector operations of process classes. See “Options on qconnector Operations” on page 77 for more information.

### **Security**

This is only relevant to client-visible records.

This option determines whether security will be applied to this operation. If security is switched on for an operation, then the generator will generate code in the RIL which checks whether the user is authorized to invoke the operation. If the user is not authorized to invoke the operation an exception will be thrown.

## SQL

This is only relevant to entity operations of stereotype ns and nsmulti.

This option allows the developer to supply the SQL code to be executed by the operation. The generator converts the supplied SQL into DAL (Java and SQL) code.

For more information about this option, see “Using Handcrafted SQL in Non-Standard Entity Operations” on page 59.

## Transactional

This allows you to specify whether a transaction is started for an operation. This is only relevant to client-visible operations.

## Where

This is only relevant to readmulti and nsreadmulti operations of entity classes.

This allows the developer to specify a customWHERE clause for the generated SQL used by the DAL code for this operation.

## Operation Parameter Options

### Mandatory Fields

This option allows the developer to specify mandatory fields for any given parameter. Mandatory fields are fields that must be populated when displayed on a client page.

The option value must be populated with a single line, comma delimited string.

Consider the following operation:

```
public interface Employer
{
    public void updateEmployerDetails(
        PersonDetails personDtls
        EmploymentDetails employmentDetails)
        throws AppException, InformationalException;
}
```

*Figure 6. Operation Signature*

The pseudo-code for the structures involved as parameters in this operation is outlined below:

```

// Note that since a person can have two addresses,
// PersonDetails aggregates AddressDetails twice
// - "homeAddress" and "workAddress".
struct PersonDetails {
    String firstName;
    String surname;
    AddressDetails homeAddress;
    AddressDetails workAddress;
}
// The role name for the struct aggregation between
// PersonDetails and AddressDetails "homeAddress" struct
// is set to "homedtIs"
struct AddressDetails {
    String addressLine1;
    String addressLine2;
    String city;
    String country;
}
// Note that EmploymentDetails aggregates AddressDetails once.
// The role name for the struct aggregation between
// EmploymentDetails and AddressDetails "employerAddress"
// struct is set to "employmentdtIs"
struct EmploymentDetails {
    String employerName;
    Date employmentStartDate;
    AddressDetails employerAddress;
}

```

Figure 7. Pseudo-Code for Parameter Structures

In this example, we want to make the following fields mandatory for our operation parameter:

For the *personDtIs* parameter:

- the person's first name; and
- the first line of the PersonDetails home address.

For the *employmentDetails* parameter:

- the first line of the employer's address.

Set theMandatory Fields option of parameter *personDtIs* to:

*firstName, homeAddress.addressLine1*

Set theMandatory Fields option of parameter *employmentDetails* to:

*employerAddress.addressLine1*

Therefore, if adding mandatory fields that are contained in structures aggregated by the parameter type class, they must be fully qualified by the relevant aggregation role names as shown above.

---

## Entity Operations Overview

A database operation is an operation of an entity class whose stereotype is recognized by the Cúram generator. The generator generates data-access Java code that is based on the stereotype for these operations.

The generator treats all other operations as if their stereotype was blank and produces Java interfaces and factories for them. The generator does not generate data-access code for these other operations - you must provide your own implementation for these operations.

## Standard Operations

### Standard Single-Record Operations

Standard single-record operations are the most basic type of operation provided in that only a single row from the database is returned and no arguments are required to be modelled as the code generator assumes standard key and details structs where appropriate.

These operations are represented by the following operation stereotypes:

- insert
- modify
- read
- remove

### Standard Multi-Record Operations

Rather than operating on a single database table row, these operations allow for processing multiple rows. In database maintenance applications, it is often necessary to return multiple records to a user interface, from which the user selects one for processing. Batch programs also frequently operate on multiple rows of a table; for example a Bank account statement printing batch program will typically operate on the accounts of every client on record.

These operations are represented by the following operation stereotype:

- readmulti

## Non-Standard Operations

### Generated SQL Operations

Non-standard generated SQL operations are similar to the standard operations except that the arguments and return type are not assumed to be standard key and standard details structs. The developer is required to specify a struct for each argument and return type.

The attributes of the argument and return structs must be subsets of the fields of the entity.

The argument structs can be user-defined structs from the input meta-model, or the generated standard structs that are not explicitly defined in the input meta-model. Using generated standard key and details structs as the parameters to non-standard operations is equivalent to simply using standard operations.

It is important to remember that since the key struct of a non-standard generated SQL operation is defined by the developer, it is possible to define a key struct which does not uniquely identify a single record. If this happens, certain operations may not behave as expected. For example, in the case of a non-standard modify operation, all records matching the key will be modified, not just the intended record.

These operations are represented by the following operation stereotypes:

- nsinsert
- nsmodify
- nsread
- nsreadmulti
- nsremove

### Handcrafted SQL Operations

Non-standard handcrafted SQL operations are the most flexible type of operation provided by the generator. They allow the developer to specify custom parameters and SQL for the operation. No parameters are generated for ns operations except those provided by the developer. All parameters provided by the developer are replicated into all the generated layers of the application.

This type of operation is intended to be used for situations where none of the other operations are suitable. This includes joins across tables and queries which count or calculate max, etc.

These operations are represented by the following operation stereotypes:

- ns
- nsmulti

### Non-Key Operations

Non-key operations operate on all rows of a database table and so would typically be used on tables containing one row.

These operations are represented by the following operation stereotypes:

- nkmodify
- nkread
- nkreadmulti
- nkremove

### Batch Operations

Batch programs, as described in *Developing batch processes*, have operations available to them that are tailored specifically to the batch environment.

These operations are represented by the following operation stereotypes:

- batchinsert
- batchmodify

---

## Entity Insert Operations

Insert operations insert, or add, a row onto a database table.

Insert operations operate on a single row at a time. There are two types of insert operation:

- insert
- nsinsert

### Standard Insert

A standard insert operation has a stereotype insert.

## Description

Standard insert operations insert a single record onto the appropriate database table using the information passed in a standard details struct. No arguments are required to be specified for these operations in the input model. Extra arguments can be specified and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

## Use

You should use a standard *insert* operation when you want to create a new record on a database table, and you want to update each attribute.

No arguments are required to be specified for this operation in the input model. Generated standard key and details structs are assumed as arguments where appropriate.

Extra arguments can be specified for this operation and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

This pattern can also be used in conjunction with the Auto ID Field sequence number generation pattern.

## Parameter and Generator Notes

Standard insert operations use the entity's details structure as an input parameter. This is automatically generated and contains all the fields of the record.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will add the standard details struct as a parameter.

## Non-standard Insert (Generated SQL)

A non-standard insert operation has a stereotype *nsinsert*.

## Description

Non-standard insert operations insert a single record onto the database table of the parent entity with information from a non-standard details struct provided by the developer.

## Use

You should use a *non-standard insert* operation when you want to create a new record on a database table, and you do not need to update each attribute. Attributes not specified in the parameter to a non-standard insert are set to null values on the database.

Non-standard insert operations are more efficient than standard inserts, because there is less I/O to the database. It is the responsibility of the application designer to decide whether the improved efficiency is worth the extra complexity of having more operations on your entities.

You might choose to use a non-standard insert where you know the database can perform the operation significantly more efficiently or where the operation will be used by a very high volume transaction.

Non-standard insert operations take a single input parameter - a structure defining the attributes to be inserted. Each attribute of this structure must match some entity attribute by name and type.

### Parameter and Generator Notes

A Warning is displayed if a non-standard operation has non-standard details parameter, which does not include fields that cannot be null. Refer to “Null Considerations” on page 59.

Fields which are not included in the details struct will not be initialized, i.e. they will be set to <null> by the DBMS.

- *Parameters* - A non-standard details struct.
- *Return Value* - None.
- *Generator action* - None.

---

## Entity Read Operations

Read operations return one or more rows from a database table, depending on the type of operation and arguments provided.

There are a number of entity read operation types:

- read
- readmulti
- nsread
- nsreadmulti
- nkread
- nkreadmulti

### Standard Read

A standard read operation has a stereotype of read.

#### Description

Standard read operations read a single record from a database table into a standard details struct, using a standard generated key struct (i.e., the primary key) as search criteria. No arguments are required to be specified for these operations in the input model. Extra arguments can be specified and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

#### Use

You should use a standard singleton *read* operation when you want to read all of the attributes of a specific database record. Standard singleton read operations use the primary key of an entity to locate the target record. You cannot create standard singleton read operations for entities that do not have primary keys. Since the primary key of an entity is unique, a standard singleton read always returns a single database record.

#### Parameter and Generator Notes

Standard singleton read operations use the entity's key and details structures as input and output parameters respectively - these are automatically generated and are not specified in the UML meta-model.

No arguments are required to be specified for these operations in the input model. Generated standard key and details structs are assumed as arguments where appropriate.

Extra arguments can be specified and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will add the standard key struct as a parameter and the standard details struct as the return value.

## Standard Readmulti

A standard readmulti operation has a stereotype readmulti.

### Description

Standard multiple read operations use an input parameter which you designate as the key structure for the operation. The return value is a structure containing a list of the entity's details structures. You must specify the first parameter, but since the return value is automatically generated, it is not specified in the UML meta-model.

### Use

You should use a standard *multiple read* operation when you want to read all of the attributes of a set of database records, based on a key that you specify. The stipulation about efficiency of keyed access, as described for non-standard read, modify and remove operations, applies equally well to multiple reads - it is up to the designer to ensure efficient use of database indices.

### Parameter and Generator Notes

A standard readmulti operation takes a partial key struct, and returns a list of standard details structs; every record matching the criteria is returned in the list.

By default, the records in a readmulti are unsorted and are returned in arbitrary order.

This can be changed by using the `order` By option of the readmulti operation. This option takes a list of the fields of the entity and sorts them in ascending order.

- *Parameters* - A key struct to specify the search criteria for which record(s) to retrieve. The members of the struct must be a subset of the standard details struct for the entity.
- *Return value* - None.
- *Generator action* - The generator will create a list wrapper for the standard details struct for the entity, and add this as the return value for the operation.

## Non-standard Read (Generated SQL)

A non-standard read operation has a stereotype nsread.

### Description

Non-standard read operations read a single record from a database table into a details struct, using an key struct as search criteria.

### Use

You should use a *non-standard read* operation when you want to read a subset of the attributes on a database record, or you want to use a key other than the primary key of the entity. Non-standard operations use a key that you specify to

locate the target record. It is not possible to guarantee at development time that only one record will be targeted. If there is more than one record in the result set, a runtime error is generated.

Non-standard read operations can be more efficient than standard ones because they result in less database I/O.

As with any operation where you specify the key yourself, there is no guarantee that the database will be able to access the target records efficiently - it is up to the designer to ensure that appropriate indices are defined to ensure this.

### Parameter and Generator Notes

Non-standard read operations use key and details structures - as input and return types respectively - that you must create yourself and specify as operation parameters in the UML meta-model. Each attribute of each of these structures must match some entity attribute by name and type. It is possible to use standard (generated) key or details structures also.

- *Parameters* - A non-standard key struct to specify the record to be retrieved. The key must be capable of uniquely identifying a single record. If more than one record matches the criteria, an exception will be thrown.
- *Return Value* - A non-standard details struct into which the data is retrieved.
- *Generator action* - None.

## Non-standard Readmulti (Generated SQL)

A non-standard readmulti operation has a stereotype nsreadmulti.

### Description

Non-standard readmulti operations take a partial key struct and a details struct as input meta-model parameters. They return a list of the provided details struct; every record matching the criteria is returned in the list.

The only difference between a non-standard readmulti and a standard readmulti is that a non-standard readmulti must specify a return value whereas for a standard readmulti this is assumed to be the standard generated details struct for the entity. For non-standard readmulti the developer is required to specify a struct as the return value of the operation. The fields of this struct must be a subset of the fields of the entity.

### Use

You should use a *non-standard multiple read* operation when you want to read a subset of the attributes of a set of database records, based on a key that you specify. It is up to the designer to ensure efficient use of database indices when reading based on this key.

### Parameter and Generator Notes

Like standard operations, non-standard multiple read operations use an input parameter which you designate as the key structure for the operation. The return value that you specify in the UML meta-model is a structure containing the attributes that you want returned for each record read from the database. The return value in the generated code is a list of the structure that you specified in the meta-model (the structure containing the list is automatically generated).

- *Parameters* - A non-standard key struct to specify the search criteria for which record(s) to retrieve. The members of the struct must be a subset of the fields of the entity.

- *Return Value* - A non-standard details struct to specify which attributes are to be returned from the readmulti operation. The members of this struct must be a subset of the fields of the entity.
- *Generator action* - The generator will create a list wrapper for the non-standard details struct specified by the developer, and use this as the return value for the operation.

## Non-key Read

A non-key read operation has a stereotype nkread.

### Description

Non-key read operations read the only record from a database table into a standard details struct.

Non-key operations, as the name suggests, do not take a key parameter. They operate by executing SQL statements which do not have a where clause; i.e. they operate on all rows on a table.

For a non-key read operation, there should be a single row on the table - this type of operation is typically used to read a value from a control table which contains a single record.

There is no such thing as a non-key insert operation since insert operations do not require a key parameter.

### Use

You should use a *non-key singleton read* operation when you want to read a record from a database table on which there is a single record. A runtime error is generated if the database table contains more than one record.

This operation type is typically used for control tables containing a single record.

### Parameter and Generator Notes

Non-key singleton read operations take no parameters and the generator automatically adds the standard details struct for the entity as the return type.

If more than one record exists on the table, an exception is thrown.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will add the standard details struct as the return value.

## Non-key Readmulti

A non-key readmulti operation has a stereotype nkreadmulti.

### Description

Non-key readmultis are very similar to standard readmulti operations, the only difference being that they return all rows of a table rather than those which match a partial key. They operate by executing SQL statements which do not have a where clause; i.e. they operate on all rows on a table.

Non-key operations, as the name suggests, do not take a key parameter. They operate by executing SQL statements which do not have a where clause; i.e. they operate on all rows on a table.

For a non-key read operation, there should be a single row on the table - this type of operation is typically used to read a value from a control table which contains a single record.

There is no such thing as a non-key insert operation since insert operations do not require a key parameter.

### Use

You should use a *non-key multiple read* operation when you want to read all of the attributes of all of the records on a database table.

### Parameter and Generator Notes

Non-key multiple read operations take no key argument. The return value is a structure containing a list of the entity's details structures as an output parameters. You specify no parameters in the UML meta-model (effectively the same interface and behavior as a standard multiple read except there is no key argument).

Generated non-key readmulti operations in the RIL and BOL have one parameter; a list details struct, i.e. a list of standard details structs for the entity.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will create a list wrapper for the standard details struct for this entity, and use this as the return value.

---

## Entity Update Operations

Update operations modify data in one or more rows of the database table, depending on the type of operation and arguments provided.

These are a number of entity update operation types:

- modify
- nsmodify
- nkmodify

### Standard Modify

A standard modify operation has a stereotype modify.

#### Description

Standard modify operations modify a specific record on an database table. No arguments are required to be specified for these operations in the input model. The record to be modified is specified using a generated standard key struct and the modified data is contained in a generated standard details struct. Extra arguments can be specified and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

#### Use

You should use a standard *modify* operation when you want to update all the attributes on a specific database record. Standard modify operations use the primary key of an entity to locate the target record. You cannot create standard modify operations for entities that do not have primary keys. Since the primary key of an entity is unique, a standard modify always updates a single database record.

The standard modify pattern can also be used in conjunction with the Optimistic Locking pattern.

### Parameter and Generator Notes

Standard modify operations use the entity's key and details structures as input parameters - these are automatically generated and are not specified in the UML meta-model.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will add the standard key struct and standard details struct as parameters.

## Non-standard Modify (Generated SQL)

A non-standard modify operation has a stereotype nsmodify.

### Description

Non-standard modify operations update records on the database table of the parent entity with information from a non-standard details struct provided by the developer.

### Use

You should use a *non-standard modify* operation when you want to update a subset of the attributes on a database record or records. Non-standard modify operations use a key that you specify to locate the target records, and this may result in multiple records being updated. You also specify which attributes of the entity are to be updated.

Non-standard modify operations can be more efficient than standard ones because they result in less database I/O, and the database may not have to update as many indices as would be the case for a standard modify operation.

### Parameter and Generator Notes

Non-standard modify operations use non-standard key and details structures as input parameters that you must create yourself and specify as operation parameters in the UML meta-model. Each attribute of each of these structures must match some entity attribute by name and type. It is possible to use standard (generated) key or details structures also.

- *Parameters* - A non-standard key struct to specify the record to be modified. Note that this non-standard key may specify multiple records. In this case, all records matching the non-standard key will be updated.

A non-standard details struct containing the updated version of the data.

- *Return Value* - None.
- *Generator action* - None.

## Non-key Modify

A non-key modify operation has a stereotype nkmodify.

### Description

Non-key modify operations modify all records on a database table with the information from a standard generated details struct.

Non-key operations, as the name suggests, do not take a key parameter. They operate by executing SQL statements which do not have a where clause; i.e. they operate on all rows on a table.

For a non-key read operation, there should be a single row on the table - this type of operation is typically used to read a value from a control table which contains a single record.

There is no such thing as a non-key insert operation since insert operations do not require a key parameter.

### Use

You should use a *non-key modify* operation when you want to update all of the records on a database table. The attribute values of each record are set to those you specify in the parameter to the non-key modify function.

Typically you would only use a non-key modify operation for control tables containing only one record.

### Parameter and Generator Notes

Non-key modify operations use the entity's details structure as an input parameter - this is automatically generated and is not specified in the UML meta-model.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - The generator will add the standard details struct as a parameter.

---

## Entity Delete Operations

Delete operations remove one or more rows from the database table, depending on the type of operation and arguments provided.

There are a number of delete operation types:

- remove
- nsremove
- nkremove

### Standard Remove

A standard remove operation has a stereotype remove.

#### Description

Standard remove operations delete a specific record from a database table. No arguments are required to be specified for these operations in the input model. The record to be deleted is specified using a generated standard key struct. Extra arguments can be specified and these arguments can be accessed by exit points for the operation, they do not have any effect on any of the generated code.

#### Use

You should use a standard *remove* operation when you want to delete a specific database record. Standard remove operations use the primary key of an entity to locate the target record. You cannot create standard remove operations for entities that do not have primary keys. Since the primary key of an entity is unique, a standard remove always deletes a single database record.

#### Parameter and Generator Notes

Standard remove operations use the entity's key structure as an input parameter - this is automatically generated and is not specified in the UML meta-model.

- *Parameters* - None.

- *Return value* - None.
- *Generator action* - The generator will add the standard key struct as a parameter.

## Non-standard Remove (Generated SQL)

A non-standard remove operation has a stereotype nsremove.

### Description

Non-standard remove operations delete records from the database table of the parent entity matching the information in a key struct provided by the developer.

### Use

You should use a *non-standard remove* operation when you want to delete a database record or records, based on a key that you specify. If the key you specify is not unique, multiple database records are deleted.

As with any operation where you specify the key yourself, there is no guarantee that the database will be able to access the target records efficiently - it is up to the designer to ensure that appropriate indices are defined to ensure this.

### Parameter and Generator Notes

Non-standard remove operations use a key structure as an input parameter that you must specify in the UML meta-model. Each attribute of this key must match some entity attribute by name and type.

**Note:** When using segmented tablespaces with DB2 for z/OS (which are the default for version 9), IBM has changed the behavior of the JDBC driver as per this Technote: <http://www-01.ibm.com/support/docview.wss?uid=swg21244002>. Therefore, a RecordNotFoundException error will not be thrown when a negative row count is returned (i.e., a DELETE FROM with no predicate).

- *Parameters* - A non-standard key struct to specify the record to be modified. Note that this non-standard key may specify multiple records. In this case, all records matching the non-standard key will be deleted.
- *Return Value* - None.
- *Generator action* - None.

## Non-key Remove

A non-key remove operation has a stereotype nkremove.

### Description

Non-key remove operations remove all the records from a database table.

Non-key operations, as the name suggests, do not take a key parameter. They operate by executing SQL statements which do not have a where clause; i.e. they operate on all rows on a table.

For a non-key read operation, there should be a single row on the table - this type of operation is typically used to read a value from a control table which contains a single record.

There is no such thing as a non-key insert operation since insert operations do not require a key parameter.

## Use

You should use a *non-key remove* operation when you want to delete all of the records from a database table.

## Parameter and Generator Notes

Non-key remove operations take no parameters.

- *Parameters* - None.
- *Return value* - None.
- *Generator action* - None.

---

## Entity Batch Operations

Batch operations insert or remove a large number of records from the database.

There are two batch operation types:

- `batchinsert`
- `batchmodify`

### BatchInsert

A batch insert operation has a stereotype `batchinsert`.

#### Description

Batch insert operations are intended to be used whenever a large amount of records are to be inserted into the database. By batching operations together, the number of round trips to the database is reduced and performance is improved.

Batch insert operations have a similar signature to non-standard insert operations and can be called in the same way. However, when a batch insert is invoked the record is not written immediately to the database. The insert statement is instead added to a batch of statements stored locally by the Cúram infrastructure by calling the `java.sql.PreparedStatement.addBatch` method. Once the batch has reached the desired size, it must be executed by calling the `$execute` method of the operation.

**Note:** The `$execute` method is never called automatically. It must be called from code written by the developer. If the entity object is destroyed without calling its `$execute` method, any pending (not executed) batched inserts will be discarded.

This means that batched inserts or modifies cannot be spread across multiple client invocations in an online environment because all entity objects are destroyed at the end of each invocation (transaction).

The `$execute` method of the operation calls the `executeBatch` method of `java.sql.PreparedStatement` and returns the result of this call which is an array of integers (`int []`). Each entry in this array corresponds to one statement in the batch and indicates how many records were affected by that statement. For example, for a successful batch of inserts, each entry of the array should be 1 to indicate that each statement caused one record to be written to the database. If one statement violated a unique constraint, its corresponding array entry would contain a zero. A returned value of `java.sql.Statement.EXECUTE_FAILED` indicates that the command failed to execute successfully.

The JDK documentation for `java.sql.PreparedStatement` provides further details regarding the information in this array, and how queued statements are executed.

The maximum number of statements in a batch is determined by the application property `curam.db.batch.limit` (default value = 30), or can be set for an individual operation by calling its `$setBatchSize(int)` method. The optimal size of a batch depends on many factors such as record size, database configuration and database vendor and can be different for each individual batch operation. It is the responsibility of the developer or DBA to determine this value.

If the batch limit is exceeded, an `AppException` (`INFRASTRUCTURE.ID_BATCH_SIZE_LIMIT_HAS_BEEN_REACHED`) is thrown by the batch insert operation. In this case the developer should simply call the `$execute` method of the operation, and then continue as before.

## Use

You should use a `batchinsert` operation when you wish to insert a large number of records to the same entity in a single transaction.

## Parameter and Generator Notes

A `batchinsert` operation takes a single input parameter - a structure defining the attributes to be inserted. Each attribute of this structure must match some entity attribute by name and type.

A warning is displayed if a `batchinsert` operation has non-standard details parameter, which does not include fields that cannot be null. Refer to section in ““Null” Considerations” on page 59.

- *Parameters* - A non-standard details struct.
- *Return Value* - None.

*Generator Action* - The generator adds the following methods to a class containing a batch insert operation:

- `public void operationName$setBatchSize(final int newBatchLimit)` - This method sets the batch limit for the operation (overrides the value of the `curam.db.batch.limit` property).
- `public int[] operationName$execute()` throws `AppException`, `InformationalException` - This method executes the currently queued batch of statements for the operation.

## BatchModify

A batch modify operation has a stereotype `batchmodify`.

### Description

Batch modify operations are similar to batch insert operations except, as the name suggests, they are used to modify existing records rather than to insert new records.

### Use

You should use a *batch modify* operation when you wish to modify a large number of records on the same entity in a single transaction.

### Parameter and Generator Notes

A *batch modify* operation uses non-standard key and details structures as input parameters that you must create yourself and specify as operation parameters in the UML meta-model. Each attribute of each of these structures must match some entity attribute by name and type. It is possible to use standard (generated) key or details structures also.

A warning is displayed if a batchmodify operation has non-standard details parameter, which does not include fields that cannot be null. Refer to section in ““Null” Considerations” on page 59.

- *Parameters* - A non-standard key struct to specify the record to be modified. Note that this non-standard key may specify multiple records. In this case, all records matching the non-standard key will be updated.

A non-standard details struct containing the updated version of the data.

- *Return Value* - None.

*Generator action* - The generator adds the following methods to a class containing a batch modify operation:

- `public void operationName$setBatchSize(final int newBatchLimit)` - This method sets the batch limit for the operation (overrides the value of the `curam.db.batch.limit` property).
- `public int[] operationName$execute() throws ApplicationException, InformationalException` - This method executes the currently queued batch of statements for the operation.

**Note:** batchmodify operations cannot be spread across multiple client-server invocations (transaction). A batchmodify operation can only be used in an online transaction if the batch is executed before the end of the transaction.

---

## Entity Handcrafted SQL Operations

Non-standard (ns) operations with handcrafted SQL can be utilized against the database.

These are two types of non-standard operation:

- ns
- nsmulti

### Non-standard

A non-standard operation has a stereotype ns.

#### Description

All parameters for a ns operation must be structs. This is because the parameters are replicated in the Data Access Layer (DAL) and the DAL allows parameters to be structs only.

The return value for a ns operation must also be a struct. Similar to parameters for ns operations, the DAL allows return values to be structs only.

The developer *must* provide SQL with all ns operations; no SQL is automatically generated.

Non-standard operations must belong to an entity class. However, the SQL query can operate on any database table, it does not have to operate on only the database table belonging to the entity class; i.e., it can be used to perform SQL joins across tables.

For details on how to specify SQL in an operation, see “Using Handcrafted SQL in Non-Standard Entity Operations” on page 59.

## Use

You should use a *non-standard* operation for a database operation which is too complex for any of the above operations and which does *not* retrieve multiple records. Examples of such operations are:

- queries whose where clause contains comparisons other than equals, such as less-than, greater-than, etc.;
- queries or commands which operate on more than one database table;
- queries which return something other than an attribute of a table, such as the results of max and count functions.

The developer must specify the SQL to be executed and can specify zero or many parameters for the operation. All parameters must be structs and must be flat, i.e. they cannot aggregate other structs.

The handcrafted SQL can perform any database operation provided that a cursor is not required. This includes single-record-reads, single or multiple record updates and deletes, and joins across multiple database tables. This is because the parameter structs cannot aggregate other structs. If your handcrafted SQL requires a cursor then an *nsmulti* operation should be used.

## Parameter and Generator Notes

- *Parameters* - Struct(s).
- *Return value* - Struct.
- *Generator action* - None.

## Non-standard multi

A non-standard multi operation has a stereotype *nsmulti*.

### Description

Non-standard multi operations are similar to non-standard operations except for the following restrictions:

- There must be either zero or one parameter;
- The operation must return a struct;
- The SQL for the operation must perform a *readmulti*.

Typically, this type of operation is intended to be used for doing *readmulti* operations which join two or more database tables.

This is the only entity operation that cannot utilize additional parameters, which is usually done to provide extra parameters to exit points in the Business Object Layer (BOL). Operations of this stereotype can have either zero or one parameter only. You cannot add any extra parameters to this operation.

## Use

You should use *non-standard multiple* operations to either:

- retrieve attributes from multiple database tables, performing a relational join across the tables, or;
- retrieve attributes from one or more database tables when the selection criteria is too complex to use a *readmulti* or *nsreadmulti*. For example, if the where clause contains comparisons other than equals such as less-than, greater-than, etc.

A non-standard multiple operation is very similar (from a modeling perspective) to a non-standard multiple read operation (*nsreadmulti*). The major difference is that

the designer must specify the SQL to be executed. This enables multiple database tables to be referenced and/or complex where clauses to be specified.

### Parameter and Generator Notes

- *Parameters* - key parameter [optional].
- *Return value* - list-details parameter.
- *Generator action* - The generator will create a list wrapper for the return-value struct specified by the developer, and use this as the return value for the operation.

You cannot specify any extra parameters for this operation.

### Example 1 - nsmulti with a Single (List) Parameter

Consider an operation in the input meta-model to list every single transaction in the system whose amount was for less than one dollar.

The following struct is defined in the model and will be used to contain the information about each transaction. The type of each attribute of the struct is not relevant here and has been omitted for clarity.

- struct class: MinorTxDetails

Attribute	Domain
txDate	DATE
txAccountNumber	ACCOUNT_NUMBER
txAmount	AMOUNT

The table below shows an entity defined in the model with some of the attributes, which will be used by the nsmulti operation `getMinorTransactions()`, returning an instance of `MinorTxDetails`.

- entity class: BankAccount

Attribute	Domain
details txDate	DATE
details txAccountNumber	ACCOUNT_NUMBER
details txAmount	AMOUNT
details txTellerNumber	TELLER_NUMBER

The SQL for the operation (which must be supplied in the model by the developer) is as follows in "Example 1 - nsmulti with a Single (List) Parameter":

```
SELECT txAccountNumber, txDate, txAmount
INTO
  :txAccountNumber,
  :txDate,
  :txAmount
FROM   BankAccount
WHERE  txAmount < 1;
```

Figure 8. SQL for nsmulti with a single (list) parameter

This is all that has to be provided by the developer, the remainder is produced by the generator and is shown below for illustrative purposes.

The following pseudo code in “Example 1 - nsmulti with a Single (List) Parameter” on page 55 describes the structs used in this operation. The actual Java structs corresponding to the structs defined in the model are produced by the code generator.

```
struct MinorTxDetails {
    txDate;
    txAccountNumber;
    txAmount;
};

// this is a generated list wrapper:
struct MinorTxDetailsList {
    sequence <MinorTxDetails> dtls;
};

// this is the standard details struct for the entity
// just to show where its attributes are kept:
struct BankAccountDtls {
    txAccountNumber;
    txDate;
    txAmount;
    txTellerNumber;
}
```

Figure 9. Pseudocode for generated structs for use by nsmulti operation

The Java interface for this entity class - complete with the nsmulti operation is produced by the code generator and would look like this:

```
public interface BankAccount {

    // This is our "nsmulti" operation. Note how the
    // generator has transformed the parameter of this function
    // from "MinorTxDetails" to a "MinorTxDetails
    // List"
    public MinorTxDetailsList getMinorTransactions()
        throws ApplicationException, InformationalException;
};
```

Figure 10. Generated Java interface for nsmulti operation

“Example 1 - nsmulti with a Single (List) Parameter” on page 55 demonstrates how the developer would write handcrafted Java code to call this method and to iterate through each element returned by the method:

```
<ProjectPackage>.intf.BankAccount bankAccount
    = <ProjectPackage>.fact.BankAccountFactory.newInstance()

double theTotalAmount = 0;

// Call the operation:
MinorTxDetailsList txList
    = bankAccount.getMinorTransactions();

// iterate through the set of results.
for (int i = 0; i < txList.dtls.size(); i++) {
    MinorTxDetails currentTx = txList.dtls.item(i);

    theTotalAmount += currentTx.txAmount;
}
```

Figure 11. Calling a nsmulti operation from handcrafted Java code (one parameter)

## Example 2 - nsmulti with Two Parameters (Key + List)

For this example, we will slightly modify the functionality of the previous example.

Instead of returning all transactions for less than one dollar, in the whole system, it will return only the transactions *for one account* which were less than one dollar.

Another parameter is required to specify the account number we are interested in. Since nsmulti is a database operation and database operations require all parameters to be structs, we must use a struct for our account number parameter even though the struct will have only one field.

Note that the account number field appears in various guises - txAccountNumber, txAccountNum, theAccountID. Unlike the other database operations, the names of attributes do not have to correspond when used in ns or nsmulti operations, the handcrafted SQL can reference the different field names as appropriate.

- struct class: AccountNoWrapper

Attribute	Domain
txAccountNumber	ACCOUNT_NUMBER

This struct can now be used as an input argument to the nsmulti operation `getMinorTransactions(theAccountID : AccountNoWrapper)`, returning an instance of `MinorTxDetails` for the entity below:

- entity class: BankAccount

Attribute	Domain
details txDate	DATE
details txAccountNumber	ACCOUNT_NUMBER
details txAmount	AMOUNT
details txTellerNumber	TELLER_NUMBER

The SQL for the operation is shown in “Example 2 - nsmulti with Two Parameters (Key + List)”:

```
SELECT txAccountNumber, txDate, txAmount
INTO
    :txAccountNumber,
    :txDate,
    :txAmount
FROM    BankAccount
WHERE   (txAmount < 1)
AND    (txAccountNumber = :txAccountNum);
```

Figure 12. SQL for nsmulti with a key and list parameters

This is all that has to be provided by the developer, the remainder is produced by the generator and is shown below for illustrative purposes.

The following pseudo-code “Example 2 - nsmulti with Two Parameters (Key + List)” describes the structs used in this operation (The actual Java structs corresponding to the structs defined in the model are produced by the code

generator.):

```
struct MinorTxDetails {
    txDate;
    txAccountNumber;
    txAmount;
};

// this is a generated list wrapper:
struct MinorTxDetailsList {
    sequence <MinorTxDetails> dtls;
};

struct AccountNoWrapper {
    txAccountNum;
}

// this is the standard details struct for the entity
// just to show where its attributes are kept:
struct BankAccountDtls {
    txAccountNumber;
    txDate;
    txAmount;
    txTellerNumber;
}
```

Figure 13. Pseudocode for generated structs for use by nsmulti with key and list parameters

The Java interface for this entity class - complete with the nsmulti operation - is produced by the code generator and would look like the following:

```
public interface BankAccount {

    // This is our "nsmulti" operation. Note how the
    // generator has transformed the return value of this
    // function from "MinorTxDetails" to a
    // "MinorTxDetails
    //     List
    //     "
    public MinorTxDetailsList getMinorTransactions
        (AccountNoWrapper theAccountID)
        throws ApplicationException, InformationalException;
};
```

Figure 14. Generated Java interface for nsmulti operation with key and list parameters

“Example 2 - nsmulti with Two Parameters (Key + List)” on page 57 demonstrates how the developer would write handcrafted Java code to call this method and to iterate through each element returned by the method:

```

<ProjectPackage>.intf.BankAccount bankAccount
    = <ProjectPackage>.fact.BankAccountFactory.newInstance();

AccountNoWrapper accNoWrapper = new AccountNoWrapper();

accNoWrapper.txAccountNum = "57033186";

double theTotalAmount = 0;

// Call the operation:
MinorTxDetailsList txList
    = bankAccount.getMinorTransactions(accNoWrapper);

// iterate through the set of results.
for (int i = 0; i < txList.dtls.size(); i++) {
    MinorTxDetails currentTx = txList.dtls.item(i);

    theTotalAmount += currentTx.txAmount;
}

```

Figure 15. Calling a *nsmulti* operation from handcrafted Java code (two parameters)

## Using Handcrafted SQL in Non-Standard Entity Operations

### Overview

For entity operations of stereotype *ns* and *nsmulti* the developer is required to specify the SQL to be used in the Cúram Data Access Layer (DAL).

These queries have access to all the tables on the database and to all the parameters of the operation.

### Using Host Variables

Host variables in SQL directly reference fields in the parameter struct or return value struct.

The rules for using host variables are as follows:

- host variables must be prefixed with a colon (:);
- host variables are case sensitive.

For example:

*:surname*

- if a field in the parameter struct or return value struct is a result of aggregation then the role name of aggregation is used for host variable.

For example:

*:dtls* (see “One-to-One Aggregation” on page 62)

### “Null” Considerations

When writing a handcrafted SQL statement, it is important to note that some Cúram datatypes are stored as null on the database if they are empty (i.e. in their initial state), so when searching for these records your query must search for “ null ” rather than an empty string. For example:

**Incorrect:**

```
SELECT ... INTO ... FROM ... WHERE someStringColumn = '';
```

**Correct:**

```
SELECT ... INTO ... FROM ...WHERE someStringColumn is null;
```

In general, if the Cúram data type corresponds to a Java class (as opposed to a primitive Java type) then its empty state is stored on the database as a null. If the data type corresponds to a primitive Java type then a null on the database is not a valid value for it and theAllow NULLs on this database field option defaults to no. If necessary this default can be overridden.

**Note:** TheAllow NULLs on this database field option controls the NOT NULL qualifier in generated DDL in an inverted way. Setting this option to no causes the NOT NULL qualifier to be added; setting it to yes causes the qualifier to be omitted.

The following table shows which Cúram data types can be represented as a null on the database.

*Table 8. Data types and nulls*

Datatype	Nulls allowed
SVR_BLOB	yes
SVR_BOOLEAN	no
SVR_CHAR	no
SVR_DATE	yes
SVR_DATETIME	yes
SVR_DOUBLE	no
SVR_FLOAT	no
SVR_INT8	no
SVR_INT16	no
SVR_INT32	no
SVR_INT64	yes
SVR_MONEY	no
SVR_STRING	yes

### For Update Considerations With DB2 for z/OS

If running against a DB2 for z/OS database, any handcrafted SQL that explicitly uses a FOR UPDATE clause may need to be modified to prevent RecordLockedException errors from being thrown. If the particular SQL statement is invoked simultaneously by multiple users, you should consider using FOR UPDATE WITH RS USE AND KEEP UPDATE LOCKS instead. The locking behavior of DB2 for z/OS is subtly different to that of DB2 on distributed platforms. The KEEP UPDATE LOCKS syntax ensures that the locking behavior with DB2 for z/OS is the same as it is on distributed platforms.

### SQL Example 1

Consider an example where the entity Employer has a method CountEmployers (stereotype ns) which returns the number of records in the Employer table.

The following struct is required to return the result, since stereotyped entity operations cannot return primitive types:

```
public final class LongWrapper
implements Serializable, DeepCloneable {
    /**
     * LONG_TYPE -> long
     */
    public long longValue = 0;
}
```

Figure 16. Struct for return result

The Java interface for this operation would look like the following extract:

```
public interface Employer
{
    public LongWrapper countEmployers()
        throws AppException, InformationalException;
}
```

Figure 17. Java Interface

Finally, the SQL to implement this query is:

```
SELECT count(*)
INTO :longValue
FROM Employer;
```

Figure 18. SQL Implementation

Note that we do not need to specify the name of the LongWrapper class, we simply reference the name of the *longValue* attribute within that class because the INTO clause is automatically assumed to reference the return value.

Thus if an attribute with the same name is used in the input parameter struct and return value struct then it is assumed that INTO clause references the attribute of return value struct.

## SQL Example 2

This example shows how to use parameter host variables and expands the previous example by adding another method which updates a numeric field on one record of the Employer table.

```
public interface Employer
{
    public void setEmployerSize(EmployerKey empKey,
        LongWrapper newSize)
        throws AppException, InformationalException;
    public LongWrapper countEmployers() throws AppException; }
}
```

Figure 19. Java Interface

The following struct is required to contain the primary key for the employer:

```

public final class EmployerKey
implements Serializable, DeepCloneable {
    /**
     * REFERENCE_NUMBER -> String
     */
    public String employerNumber = "";
}

```

Figure 20. Struct for employer key

The SQL statement for this method is:

```

UPDATE Employer
SET size = :2.longValue
WHERE employerNumber = :employerNumber;

```

Figure 21. SQL Implementation

Note that since *longValue* is contained in the second parameter it is necessary to qualify it with 2.. Unqualified parameter references are assumed to reference the first parameter.

The SQL statement below qualifies both parameters and is equivalent to the one above:

```

UPDATE Employer
SET size = :2.newSize.longValue
WHERE employerNumber = :1.employerNumber;

```

Figure 22. SQL Implementation with qualified parameters

---

## Aggregation

Aggregation allows you to embed or nest instances of one type of class within another type of class.

The Cúram generator supports two types of aggregation relationships: one-to-one and one-to-many. One-to-one aggregation embeds a single instance of one class within another. One-to-many aggregation embeds a sequence of one class within many others.

The main use for aggregation in the generator is to represent sequences in the input meta-model.

The generator permits the following aggregation configurations:

- structs can aggregate structs;
- structs can aggregate entities.

### A Special Case

The generator supports the aggregating of standard details structs, even though they do not appear in the input model. Standard details structs are aggregated by aggregating the entity class which “owns” the standard details struct.

### One-to-One Aggregation

The following example describes how to aggregate a struct class, *PersonDetails*, into to another struct class, *PersonDetailsWrapper*, using one-to-one aggregation.

To create a one-to-one aggregation create an Rational Software Architect diagram and do the following:

- Add classes `PersonDetails` and `PersonDetailsWrapper` to the diagram;
- In the diagram drag the appropriate arrowhead (appears when the mouse cursor is over the class) between the two classes with `PersonDetailsWrapper` set as the source and `PersonDetails` the target;
- Select Create Aggregation from the popup menu;
- With the aggregation relationship selected in the diagram open the General Properties tab.

This creates the aggregation relationship whereby one role corresponds to class `PersonDetailsWrapper` and the other to class `PersonDetails`. A UML role is essentially one end of a UML relationship so each relationship has two roles whose names are Role A and Role B. Exactly one of these roles - usually Role A - will have its Aggregate option set. The assignment of Role A and Role B is arbitrary. The key thing to remember is that the role which has the Aggregate box checked denotes the outermost class of the pair.

With the relationship line selected in the diagram the General Properties tab should show `PersonDetailsWrapper` in the graphic at the top of the properties sheet with the diamond associated with it. Set the following properties of the aggregation:

- The Label is optional;
- For `PersonDetailsWrapper`:
  - The Aggregation radio button should indicate Composite;
  - Multiplicity should be set to 1;
- For `PersonDetails`:
  - The Aggregation radio button should indicate None;
  - By default the role is set to "dtls";
  - Multiplicity should be set to 1 (to signify a one-to-one aggregation).

The class diagram would appear in the Rational Software Architect showing the two classes joined by the UML aggregation relationship line (diamond end touching `PersonDetailsWrapper`) and each side of the relationship showing multiplicity of one and the `PersonDetails` showing a role name of - dtls.

**Note:** The position of the diamond in the model diagram is important as it denotes the outermost class in the pair.

The generated Java code resulting from this construct would take the following format:

```
public final class PersonDetails implements
java.io.Serializable, curam.util.type.DeepCloneable {
    public String personRefNo = "";
    public String firstName = "";
};
public final class PersonDetailsWrapper implements
java.io.Serializable, curam.util.type.DeepCloneable {
    // This class has a single instance of
    // class "PersonDetails" embedded in it. PersonDetails dtls =
    // new PersonDetails();
};
```

## One-to-Many Aggregation

In this example a one-to-many aggregation is modeled, meaning that a list of one class type is embedded into the other class. Here we create `PersonDetailsList`, which aggregates a list of `PersonDetails`. To create a one-to-many aggregation, open an Rational Software Architect diagram and do the following:

- Add classes `PersonDetails` and `PersonDetailsList` to the diagram;
- In the diagram drag the appropriate arrowhead (it appears when the mouse cursor is hovering over the class) between the two classes with `PersonDetailsList` as the source and `PersonDetails` as the target;
- Select Create Aggregation from the popup menu;
- With the aggregation relationship selected in the diagram open the General Properties tab.

This creates the aggregation relationship whereby one role corresponds to class `PersonDetailsList` and the other to class `PersonDetails`.

**Note:** The position of the diamond is important as it denotes the outermost class in the pair.

With the relationship line selected in the diagram the General Properties tab should show `PersonDetailsList` in the graphic at the top of the properties sheet with the diamond associated with it.

Set the following properties of the aggregation:

- The Label is optional;
- For `PersonDetailsList`:
  - The Aggregation radio button should indicate Composite;
  - Multiplicity should be set to \*;
- For `PersonDetails`:
  - The Aggregation radio button should indicate None;
  - By default the role is set to "dtls";
  - Multiplicity should be set to 1..\* (to signify a one-to-many aggregation).

The class diagram would appear in the Rational Software Architect showing the two classes joined by the UML aggregation relationship line (diamond end touching `PersonDetailsList`) and the aggregates side of the relationship showing a multiplicity of \* and `PersonDetails` showing a multiplicity of 1..\* and a role name of - dtls.

The pseudo-code resulting from this construct would take the following format:

```

struct PersonDetails implements
java.io.Serializable, curam.util.type.DeepCloneable {

    String personRefNo = "";
    String firstName = "";
};

struct PersonDetailsList implements
java.io.Serializable, curam.util.type.DeepCloneable {

    public static class List_dtls
    extends curam.util.type.ValueList {
        public void addRef(PersonDetails s) {
            add(s);
        }
        public PersonDetails item(final int indx) {
            return (PersonDetails) get(indx);
        }
        public PersonDetails[] items() {
            PersonDetails[] result = new PersonDetails[size()];
            toArray(result);
            return result;
        }
    }

    // This class contains an embedded list of "PersonDetails":
    public final List_dtls dtls = new List_dtls();
}

```

The resulting generated struct class for `PersonDetailsList` has a field named `dtls` which provides functionality required for lists such as adding items, getting an item by index and getting the list contents as an array.

---

## Assignable

You can use generated struct classes to automatically assign values between matching fields in another struct as provided by the generated struct's super class `curam.util.type.struct.Struct`.

Consider an example of a struct, `BankBranchStruct` with several attributes:

- *bankBranchID*
- *bankId*
- *bankName*
- *bankSortCode*
- *name*
- etc.

A `BankBranchListDetails` struct class has a subset of attributes that shared with the `BankBranchStruct` class:

- *bankBranchID*
- *bankSortCode*
- *name*

The following Java code illustrates how to create these objects.

```

BankBranchStruct bankBranchStruct
    = new BankBranchStruct();
BankBranchListDetails bankBranchListDetails
    = new BankBranchListDetails();

```

Typically, the assignment from one struct to the other look like this:

```

bankBranchListDetails.bankBranchID
    = bankBranchStruct.bankBranchID;
bankBranchListDetails.bankSortCode
    = bankBranchStruct.bankSortCode;
bankBranchListDetails.name = bankBranchStruct.name;

```

The above code can be simplified by using the assign function, which becomes more significant as the size of the structs increases:

```

bankBranchListDetails.assign(bankBranchStruct);

```

An assignable relationship is one that allows further control of the specifics of the automatic assignment with the assign function. It is required when you want to do explicit field assignment between fields with differing names or to suppress the default assignment between fields of the same name.

## Explicit Field Assignment

An explicit field assignment is one where fields with different names are matched. It is represented in the model by adding an assignable relationship between the two classes, and then adding attributes to be matched to the both sides of the assignment. Any fields which are not explicitly linked will be treated as default assignment fields.

The following classes are used to illustrate this.

- entity class: Address

Attribute
addressID
addressLine1
addressLine2
addressLine3
addressLine4
cityCode
countryCode
postalCode
regionCode
comments

- struct class: BankBranchStruct

Attribute
bankBranchID
bankID
bankName
addressID

Attribute
addressLine1
addressLine2
addressLine3
addressLine4
countryCode
postalCode
regionCode
addressVersionNo
cityID

In an assignable relationship between the two classes `Address` and `BankBranchStruct` fields can be explicitly mapped; e.g. `BankBranchStruct.cityID` matched with `Address.cityCode`. In Rational Software Architect this is shown in **Role:** fields (RoleA & RoleB) of the **General** tab of the assignable relationship with the linked pair, `cityID` in one Role field and `cityCode` in the other. All the other common fields (e.g. `AddressLine1`, etc.) are handled automatically by the generator.

For instance, the generated code without the explicit field assignment would appear as shown below:

```
public curam.util.testmodel.struct.BankBranchStruct
    assign(final curam.util.testmodel.struct.AddressDtls v)
{
    addressID = v.addressID;
    addressLine1 = v.addressLine1;
    addressLine2 = v.addressLine2;
    addressLine3 = v.addressLine3;
    addressLine4 = v.addressLine4;
    countryCode = v.countryCode;
    postalCode = v.postalCode;
    regionCode = v.regionCode;
    return this;
}
```

With the explicit field assignment the following code is then added to the assign method: `cityID = v.cityCode`. The handcrafted Java to assign these structures would be as follows:

```
BankBranchStruct dtls = new BankBranchStruct();
AddressDtls addressDtls = new AddressDtls();
dtls.addressLine1 = addressDtls.addressLine1;
dtls.addressLine2 = addressDtls.addressLine2;
dtls.addressLine3 = addressDtls.addressLine3;
dtls.addressLine4 = addressDtls.addressLine4;
dtls.cityID = addressDtls.cityCode;
dtls.countryCode = addressDtls.countryCode;
dtls.postalCode = addressDtls.postalCode;
dtls.regionCode = addressDtls.regionCode;
```

By using the generated assignment operator, these lines of code can be reduced to just one line as follows:

```
bankDtls.assign(addressDtls);
```

## Suppressing Default Assignment Fields

In some situations you may not want a pair of similarly named fields to be matched. You can cause a pair of fields to be omitted from an assignment by listing one of the fields at one end of the relationship.

For the following two classes below, `PersonInfo` and `AccountInfo`, having an struct relationship, the same named fields are matched.

- struct class: `AccountInfo`

Attribute
Id
Surname
FirstName
Balance

- struct class: `PersonInfo`

Attribute
Id
Surname
FirstName

For this example we first create the objects for the `PersonInfo` and `AccountInfo` classes as described above:

```
AccountInfo account = new AccountInfo();
PersonInfo person = new PersonInfo();
```

This assignment:

```
account.assign(person);
```

is equivalent to the following three statements:

```
account.Id = person.Id;
account.Surname = person.Surname;
account.FirstName = person.FirstName;
```

By adding *Id* as a key to one end of the relationship, it is excluded from the generated assignment and now this assignment:

```
account.assign(person);
```

is equivalent to the following *two* statements; that is, the *Id* assignment will no longer be made:

```
account.Surname = person.Surname;
account.FirstName = person.FirstName;
```

## Combining structs

Sometimes you may need to populate one struct with the contents of two or more other structs.

A typical piece of Java code would look like the following:

```
BankBranchStruct dtls = new BankBranchStruct();
AddressDtls      addressDtls = new AddressDtls();
BankBranchDtls  bankBranchDtls = new BankBranchDtls();

// Copy from the "AddressDtls" struct
dtls.addressLine1 = addressDtls.addressLine1;
dtls.addressLine2 = addressDtls.addressLine2;
dtls.addressLine3 = addressDtls.addressLine3;
dtls.addressLine4 = addressDtls.addressLine4;
dtls.cityCode     = addressDtls.cityCode;
dtls.countryCode  = addressDtls.countryCode;
dtls.postalCode   = addressDtls.postalCode;
dtls.regionCode   = addressDtls.regionCode;
dtls.
    addressVersionNo
    = addressDtls.
    versionNo
    ;

// Copy from the "BankBranchDtls" struct
dtls.bankBranchID = bankBranchDtls.bankBranchID;
dtls.bankID       = bankBranchDtls.bankID;
dtls.bankSortCode = bankBranchDtls.bankSortCode;
dtls.name         = bankBranchDtls.name;
dtls.versionNo    = bankBranchDtls.versionNo;
```

Figure 23. Example Java code for combining structs

By explicitly mapping the *BankBranchStruct.addressVersionNo* attribute to the *Address.versionNo* in the assignable relationship the Java can now be written as:

```
// Copy from the "AddressDtls" struct
dtls.assign(addressDtls);

// Copy from the "BankBranchDtls" struct
dtls.assign(bankBranchDtls);
```

Figure 24. Equivalent Java code for combining structs

Note that in this case, the second assign does not overwrite the first as it happens to reference a different subset of fields, so the net effect is that the two struct contents are merged.

---

## Foreign Keys

You can create foreign keys between database tables with the Cúram generator.

A foreign key relationship between two database tables is specified in the input model by adding a relationship of stereotype *foreignkey* (one word, no spaces) between two entity classes. Optionally you can give the relationship a name, this name is then applied to the foreign key constraint added to the database. Otherwise the database chooses its own name for the constraint.

Rules when using foreign keys:

- Foreign key relationships are allowed on entity classes only.
- Fields referenced by a foreign key will be set to unique, as this is required by some databases.

- If the foreign key references the primary key of another entity, a redundant unique clause will not be generated by the generator, as the primary key is already unique.
- Foreign keys cannot be specified on subclass entities. The relationship should be specified using the actual base entity classes themselves.

## How to Add a Foreign Key to a Database Table

A foreign key is specified between a pair of entities by adding a relationship between the two classes and adding key/qualifiers to the role touching the *referenced* class. On a class diagram, this results in a line between two classes, with a box containing the key/qualifiers at the *referencing* class.

The notation for linking pairs of fields in two different classes is the same for foreign keys as for generated assignments. The class diagram will show two classes joined by a line with pairs of linked attributes in a box at one end of the line. The first name in the pair refers to an attribute in the nearer class, the second name refers to an attribute in the other class.

## Naming Primary and Foreign Key Constraints

It is possible to include a constraint name for foreign key constraints in Cúram models. The name given in the model to the foreignkey relationship will be applied to the foreign key constraint itself. If necessary this feature can be suppressed by specifying '-nonamedforeignkeyconstraint' on the generator command line.

Primary key constraints are also given names in the database. The name of each constraint is the same as that of its corresponding entity. This also results in an accompanying index of the same name. This feature can be suppressed by specifying'-nonamedprimarykeyconstraint' on the generator command line.

## Example

Consider two entity classes, BankAccount and BankTransaction, where BankAccount.*accountNo* is a foreignkey on BankTransaction. That is, the BankTransaction table (txAccountNo) must have a record on the BankAccount table with a matching accountNo value.

The tables below illustrate these two classes where the foreignkey would be between their key attributes:

- entity class: BankAccount

Attribute	Domain
key <i>accountNo</i>	ACCOUNT_NO
details clientID	CLIENT_ID
details branchLocation	BRANCH_LOCATION
details currentBalance	MONEY
details lastTransaction	DATE_TYPE
details lastStatement	DATE_TYPE

- entity class: BankTransaction

Attribute	Domain
key <i>txAccountNo</i>	ACCOUNT_NO

Attribute	Domain
details txID	TX_ID
details transactionDate	TX_DATE
details transactionType	TX_TYPE
details transactionAmount	TX_AMOUNT

This foreign key results the following DDL being generated (Oracle SQL shown):

```
ALTER TABLE BankTransaction ADD(
  FOREIGN KEY (txAccountNo)
  REFERENCES BankAccount(accountNo));
```

---

## Indexes

You can create indexes other than the primary index on database tables using the Cúram generator. Any number of indexes can be created on each table, with the usual speed versus database size trade-offs that are associated with indexes.

An index for a database table is specified in the input model by adding a relationship of stereotype index between an entity class and a struct class.

The fact that a struct is being used to represent an index does not have any side-effects on the struct apart from those mentioned. The struct can still be used as an argument to an operation. Typically the struct would be used as both a key parameter and as an index to support database accesses via this key.

Rules when using indexes:

- The relationship must be given a name. This is the name for the to the database index. The name of the struct in the relationship does not affect the index.
- The names of the attributes of the struct class must be a subset of the names of the attributes of the entity.
- The struct class must not aggregate any other classes.
- Index names must be unique within the entire model.

### How to Add an Index to a Database Table

Create a struct class whose fields are a subset of the fields of the entity class.

Add a relationship of stereotype index between the entity class and the struct. The direction of the relationship is not important.

Set the relationship name to the name which you want given to the database index.

### Naming Indexes

Developers will never explicitly reference an index but the DBA<sup>3</sup> will, so it is recommended that index names be kept as meaningful and descriptive as possible.

---

3. Data Base Administrator.

## Example

Consider the following two classes with an index relationship named BankClientMNIndex:

- entity class: BankClient

Attribute	Domain
key clientID	CLIENT_ID
details firstName	CUSTOMER_NAME
details middleName	CUSTOMER_NAME
details lastName	CUSTOMER_NAME
details address1	ADDRESS_LINE
details address2	ADDRESS_LINE
details address3	ADDRESS_LINE
details address4	ADDRESS_LINE

- struct class: MiddleNameWrapper

Attribute	Domain
middleName	CUSTOMER_NAME

The above index results in the following DDL being produced by the generator:

```
CREATE INDEX BankClientMNIndex
ON BankClient(middleName);
```

---

## Unique Indexes

A unique index is modeled by adding a uniqueindex stereotype relationship between an entity class and a struct class. The rules for modeling a unique index are the same as those for modeling a non-unique index.

Specifying a unique index for an entity causes the necessary information to be included in the generated file <Application-name>\_unique\_constraints.xml which must then be referenced from the data manager configuration file (datamanager\_config.xml).

Note that the file <Application-name>\_unique\_constraints.xml contains two sets of information:

1. Unique indexes correspond to explicit 'uniqueindex' relationships in the model and result in DDL of the form:

```
CREATE UNIQUE INDEX <index-name>
ON....
```

where ' <index-name> ' is the name of the relationship in the model.

2. Unique constraints are implicit unique constraints which are produced automatically by the generator and which are applied to all fields which are referenced by a foreign key. They correspond to foreignkey relationships in the model and result in DDL of the form:

```
ALTER TABLE <table-name> ADD
    UNIQUE...
```

or, if there is a uniqueindex for the fields referenced by the foreign key:

```
ALTER TABLE
    <table-name> ADD CONSTRAINT <constraint-name>
    UNIQUE...
```

where ' <constraint-name> ' is the name of the corresponding uniqueindex relationship in the model.

When the data manager is run, the explicit unique indexes are created before the implicit unique constraints. This allows the database to use the developer-specified unique indexes to enforce uniqueness rather than having to create and use its own system-named indexes.

For example, you might want to model a specifically named unique index to correspond to a particular foreign key in the model. The generator automatically gives the unique constraint the same name as the corresponding unique index.

---

## Generated Class Hierarchy

There is a hierarchy of classes that are generated by the server code generator that corresponds to the classes designed in the application model.

All classes are defined in the Cúram model using UML notation. A single process, facade, or entity class can contain a mixture of automatically generated and handcrafted methods. It is not desirable to store handcrafted code and generated code in the same file due to the risk of the generator overwriting handcrafted code, or vice versa.

Developer provided code is stored in a single class. Generated code is generated into a number of other classes, and the set is linked together into a hierarchy by inheritance and implementation.

Note that since struct classes do not contain operations there are no issues of separating handcrafted and generated code. Each struct class in the model corresponds to one generated Java struct class.

### Basic Hierarchy Example

This section describes the elements of the generated and required handcrafted class hierarchy for a basic entity class named `MyClass`, which does not make use of inheritance or code packages.

The UML representation of the generated Java classes of the entity class `MyClass` would show the following four classes:

- `<PackageName>.intf.MyClass`
- `<PackageName>.base.MyClass`
  - Implements, or realizes, the `intf` class.
  - It is the super class.
- `<PackageName>.impl.MyClass`
  - A subclass of the base class.
  - Contains any required (non-generated) handcrafted methods.

- `<PackageName>.fact.MyClassFactory`
  - A subclass of the `impl` class.
  - Returns an instance of the `intf` class.

Thus, there are four Java classes corresponding to the entity class in the UML model. Three of the classes have the same name as the class in the model, the fourth has the same name with the word `Factory` appended.

A further description of the classes are as follows:

1. `<ProjectName>.intf.MyClass`

This is a generated Java interface class containing all the public methods for the class.

The other classes in the hierarchy - either generated or handcrafted - will be required to provide implementations for these methods.

2. `<ProjectName>.base.MyClass`

This is a generated abstract Java class which implements the interface contained in the `intf` version of the file. It contains the following:

- The implementations of data access methods (i.e. stereotyped methods of entity classes) and connector methods.
- Abstract method declarations for exit point methods.

This is to ensure that the developer is forced to provide implementations for the exit points.

- Abstract method declarations for methods declared protected in the model.

This is to ensure that the developer is forced to provide implementations for these methods without having to expose them in the interface (`intf` layer) for the class.

3. `<ProjectName>.impl.MyClass`

This class is supplied by the developer and always inherits from the corresponding base version.

It should be declared abstract to ensure that the class cannot be instantiated directly - the class should only be instantiated using the factory mechanism. (See below.)

In this class the developer must provide implementations for all the methods declared in the class in the model for which an implementation was not produced by the generator.

While this class inherits from a generated class, it contains only handcrafted code and *no* generated code. This is so that there is no risk of developer code overwriting generated code, or generated code overwriting developer code.

4. `<ProjectName>.fact.MyClassFactory`

This is a generated Java class containing one static method: `newInstance()`. This method creates instances of the class and is the only means by which entity, facade and process classes should be instantiated.

Since a factory creates all instances of objects, it can also be used to:

- transparently create and return a customized version of the class requested. See “Replacing the Superclass” on page 87. Pre-existing code which used the original version of the class does not need to be changed.
- transparently create and return a proxy class of the requested class. The proxy class wraps the requested class (using the Java 1.3 Dynamic Proxy mechanism) and captures detailed tracing information for all interactions with the class.

The following code sample shows how an instance of MyClass is created. Note that the return type of MyClassFactory.newInstance is sample.intf.MyClass.

```
// Use the factory to create an instance:
sample.intf.MyClass myObject =
    sample.fact.MyClassFactory.newInstance();
```

Figure 25. Using a factory to create an instance of MyClass

## Hierarchy for Subclasses

This section describes the elements of the generated/handcrafted class hierarchy for a basic entity class named SubClass that inherits from MyClass.

The UML representation of the generated Java classes for the entity class SubClass would show the following four classes:

- <PackageName>.intf.SubClass
  - It inherits from the MyClass intf class.
- <PackageName>.base.SubClass
  - Implements, or realizes, the intf class.
  - It is the super class.
  - It inherits from the MyClass impl class.
- <PackageName>.impl.SubClass
  - A subclass of the base class.
  - Contains any required (non-generated) handcrafted methods.
- <PackageName>.fact.SubClassFactory
  - A subclass of the impl class.
  - Returns an instance of the intf class.

As with the previous example there are four Java classes corresponding to class SubClass. However the fact that SubClass inherits from MyClass results in two additional relationships, highlighted here:

1. Interface SubClass inherits from interface MyClass thereby ensuring that SubClass must implement all of its own declared methods plus those declared in MyClass.
2. Generated class <ProjectName>.base.SubClass inherits from handcrafted class <ProjectName>.impl.MyClass. This means that SubClass inherits the implementations of the methods from SubClass as well as their declarations, so these methods are available to SubClass and do not have to be re-implemented.

## Hierarchy for Abstract Classes

In a Cúram model the developer can mark classes abstract (See “Options” on page 18) meaning that they cannot be instantiated.

From the above example, if MyClass were qualified abstract, the following hierarchy would result:

- <PackageName>.intf.MyClass
- <PackageName>.base.MyClass
  - Implements, or realizes, the intf class.
  - It is the super class.
- <PackageName>.impl.MyClass
  - A subclass of the base class.

- Contains any required (non-generated) handcrafted methods.

The hierarchy is the same as for non-abstract classes except that no factory is generated.

## Considerations

### Access Control - private/protected/public/package

The Java language supports four levels of access control for methods and member variables. In Cúram models this is simplified to two levels: public and protected. Since the generated class hierarchy includes different classes in different packages it does not make sense to use the private and package access levels. Note that this applies only to operations in Cúram models - developers are still free to use private and public access in handcrafted Java code as they see fit.

### The Meaning of super

In Java, the super keyword is a reference to the superclass i.e. the class from which the current class (this) inherits.

In the example shown in “Hierarchy for Subclasses” on page 75 the superclass of SubClass is MyClass. However when writing handcrafted Java code for `<ProjectName>.impl.SubClass` it is important to remember that the superclass of this class is actually `<ProjectName>.base.SubClass` rather than any version of `MyClass`.

### Enforcing the Factory Mechanism

For reasons mentioned above, entity, facade and process objects should be created only by using their associated factory classes. Developers should not bypass this mechanism by using the new keyword to instantiate these classes. This can and should be enforced by making all implementation classes (i.e. all classes in the `impl` packages) abstract. Failure to make these classes abstract means that there is a risk of developers instantiating them directly with the result that class replacement will not work as expected.

## Summary

Certain individual objects in a Cúram application model appear as multiple classes in the output code. The objective of the generated class hierarchy is to ensure the following:

- The developer provides all handcrafted implementation within a single Java class.
- The public parts of the object's interface are accessible to other objects and the non-public parts of the object's interface are not accessible to other objects.
- The developer is forced to implement all of the declared interface, both public and non-public - unless the generator produces the necessary implementation.
- Objects can be subclassed and a subclass can be defined to replace its superclass transparently.
- The run time type of an object is determined by a factory, to support replacement and tracing.

---

## Cúram JMS Queue Connectors

Connectors provide a way for a Cúram application to connect to other systems by means of JMS queues.

For facade and process class operations with a stereotype of `qconnector` the generator produces code that converts the operation parameter into a JMS message, places the message on a queue, and optionally waits for another message. The message is then converted back into a `Cúram` struct and returned to the caller.

For many operations, queue connectors can be implemented without writing any handcrafted code. It is also possible to customize connectors with the use of handcrafted code. Do this if:

- The default encoding of a datatype is not suitable for your purpose. For example, you want to encode dates in the form `DD-MMM-YYYY` instead of the default format of `YYYYMMDD`.
- Your parameter struct is “complex”. For example, it contains a variable length field, or aggregates another struct.

## Domain Definition Overview

IBM `Cúram` Social Program Management connectors provide a way for a `Cúram` application to connect to other systems by means of JMS queues. For facade and process class operations with a stereotype of `qconnector` the generator will produce code that converts the operation parameter into a JMS message, places the message on a queue, and optionally waits for another message in response which is then converted back into a `Cúram` struct and returned to the caller.

For many operations, queue connectors can be implemented without writing any handcrafted code. It is also possible to customize connectors with the use of handcrafted code. You may wish to do this if:

- the default encoding of a datatype is not suitable for your purpose. For example, you may wish to encode dates in the form `DD-MMM-YYYY` instead of the default format of `YYYYMMDD`.
- your parameter struct is “complex”. For example, it may contain a variable length field, or may aggregate another struct.

## How It Works / What It Does

Connections are not created directly, but are built using a connection factory. Factory objects are stored in a JNDI namespace, insulating the JMS application from provider-specific information.

The fields in the parameter struct are scanned using Java reflection, and each is converted into a fixed length string based on its datatype. The strings are concatenated together into a JMS `BytesMessage` which is then placed on a JMS queue.

If a return type has been specified for the operation, the `Cúram` application will wait for a response message, typically on another queue. The remote system must create a correctly formatted response message and send it to the `Cúram` application within the specified timeout period. When the message is received, it is converted into an instance of the return type struct which is then returned to the caller.

## Options on `qconnector` Operations

The following options are available on `qconnector` operations:

- **JNDI name of the `QueueConnectionFactory` class**  
*Mandatory.* This specifies the name of the `QueueConnectionFactory` class in the JNDI namespace.

Queue connections are not instantiated directly but are instead created by connection factories. The connection factories are stored in the JNDI namespace of the application server.

- **JNDI name of the transmission queue**

*Mandatory.* This specifies the JNDI name of the queue onto which outgoing messages are placed.

- **Response message timeout (seconds)**

This is only relevant for operations that have a return value. The return value is obtained by receiving a response JMS message from the recipient and this timeout value is used to ensure that the application does not wait indefinitely for the response.

Default value: 30 seconds.

- **JNDI name of the reply queue**

This is only relevant for operations that have a return value. Specifies the JNDI name of the queue from which the response message should be taken.

- **Message Type**

BytesMessage or TextMessage. This allows you to specify whether a JMS BytesMessage or TextMessage is sent/received by the connector. By default the JMS connectors send and receive a JMS BytesMessage containing the bytes of a string representation of the struct parameters. If the system(s) being communicated with use a different character encoding, then these bytes may not be correctly translated by the other systems. In this case - provided the message doesn't contain any binary data - a JMS TextMessage can be used to ensure that the message is correctly translated by the other systems.

- **BytesMessage encoding character set**

Specifies the name of the character encoding to use when converting the string representation of a struct to a JMS BytesMessage, and vice versa. If this option is not specified then the default local system character encoding is used. (Usually 'Cp1252' for Microsoft Windows, 'Cp1046' for EBCDIC on IBM z/OS®, etc.) This enables you to ensure that the character encoding used for the message matches the character encoding of the other system being communicated with.

This option is not relevant if TextMessage is used.

## How to Use qconnector Operations

The following section explains how qconnector operations are represented in the meta-model and implemented on the remote system.

### Decide on Format of Message and Create the struct(s) to Correspond to the Message

The Cúram developer and the remote application developer need to agree on the format(s) of the messages passed between the two systems. This involves:

- **The format of each field in the message**

The default encoding method can be used for each field; but, see “Encoding Methods for Fundamental Types” on page 79 and “Using Customized Encoding/Decoding Classes” on page 80 for how a custom encoding methodology may be implemented.

- **The length of each field in the message**

Like the encoding, the encoded length of each field depends on the type of the field and - for some datatypes - its length as specified in the model. See “Encoding Methods for Fundamental Types” on page 79 for information on lengths of datatypes. The length of the field can be changed by implementing a custom mapper for the field.

- **The ordering of the fields in the message**

The fields appear in the message in the same order as they appear in the struct in the meta-model. The tool-bar contains a facility for changing the order of struct attributes if required.

**Add the operation to the application meta-model.**

A qconnector operation is modelled like any other process or facade class operation subject to the restrictions listed in “Rules / Restrictions.” It is also necessary to use some of the operations listed in “Options on qconnector Operations” on page 77 to specify the queue(s) and some queuing parameters.

In summary, the method should have one struct parameter, may return void or a struct, and should have options set to identify the MQSeries queues to use.

**Configure the Queues in the Application Server**

The queue connection factory, and references to the queues themselves are stored in the JNDI namespace. These JNDI names are be mapped to actual connection factories and queues in the application server configuration.

**Implement the message recipient in the remote system**

The message recipient can be any system which has access to the MQSeries queues. Typically this will be a legacy system to which access is required by the Cúram application. The target system can be either a JMS application or a basic MQSeries application.

If no response is required from the remote system, the remote system simply collects and decodes the received message, and uses it as required.

If a response message is required, i.e. if a return type has been specified for the operation, then the remote system must create a response message and send it back to the waiting Cúram application. The response message is associated with the original message using its *CorrelationID*, i.e., *the message recipient must set the CorrelationID of the response message equal to the MessageID of the original message.*

**Rules / Restrictions**

- The qconnector operation stereotype is valid in process or facade classes only.
- Connector operations must have exactly one struct parameter.
- Connector operations may have a return type of void or a struct.
- The parameter and return structs may take any form, however the generated code is only capable of mapping structs which are “flat” - structs that do not aggregate other structs - and which have only fixed length fields. For complex structs, it is necessary to implement a mapper class to map the struct to and from messages. Examples of coding and decoding complex structs are provided below.

**Encoding Methods for Fundamental Types**

*Table 9. Encoding methods*

Datatype	Encoded Width	Encoding method
SVR_BLOB	Variable	Converted directly to a padded string
SVR_BOOLEAN	1	false = 0, true = 1

Table 9. Encoding methods (continued)

Datatype	Encoded Width	Encoding method
SVR_CHAR	1	Converted directly to a 1 character string
SVR_DATE	8	yyyyMMdd
SVR_DATETIME	15	yyyyMMddThhmmss (ISO 8601 standard)
SVR_DOUBLE	25	Numeric
SVR_FLOAT	16	Numeric
SVR_INT8	1	Numeric
SVR_INT16	6	Numeric
SVR_INT32	11	Numeric
SVR_INT64	21	Numeric
SVR_MONEY	25	Numeric
SVR_STRING	Variable	converted directly to a padded string
SVR_UNBOUNDED_STRING	N/A	Not natively supported

- SVR\_BLOB and SVR\_STRING are variable in that the length of the encoded message is equal to the length specified for that type in the model. If the data in the string is less than the maximum amount allowed, space padding is appended to the data in the message to bring it up to the maximum size.
- SVR\_UNBOUNDED\_STRING is not natively supported because the string length is not known at generate time and is required for creating fixed length messages. However it is possible for the developer to implement a custom mapper to handle unbounded strings.
- Numeric datatypes are converted to right-justified human-readable strings. For example: 45678, -23123, 1000003.14159, 1.4E-45.

## Using Customized Encoding/Decoding Classes

By default the encoding method used for each field in a struct used or returned by connector operations is based on the type of the field. For example, the mapper class for `curam.util.type.DateTime` is `curam.util.connectors.mqseries.MQFieldMapper.DateTimeMapper`; for boolean fields it is `curam.util.connectors.mqseries.MQFieldMapper.BooleanMapper`.

For any individual field in any operation it is possible to override this default and specify the name of the class which should be used to map the data. Names of the custom mapper classes are specified in the properties file `QueueConnectorFieldMappers.properties` which must be included in the application classpath.

Entries in the properties file take the following format:

```
[class].[operation].[param].[field]=[mapper]
```

where

- [class] is the name of the process or facade class containing the connector operation;
- [operation] is the name of the connector operation;

- [param] is the name of the parameter - or the property return to specify the return value for the operation;
- [field] is the name of the field within the parameter struct;
- [mapper] is the fully qualified class name of the required mapper class. This must be a subclass of `curam.util.connectors.mqseries.MQFieldMapper`.

```
MyBPO.connectorOp1.dtls.phoneNumber=com.acme.util.PNMapper
MyBPO.connectorOp1.return.phoneNumber=com.acme.util.PNMapper
```

Figure 26. Sample `QueueConnectorFieldMappers.properties`

## Example 1 - Working with Variable Length Fields

In the following example a custom field mapper class is used to implement a primitive variable length field message. The variable length field is encoded by prefixing the data with a six character string containing a number which specifies the length of the data in the remainder of the string.

Note that this example only shows the implementation at the Cúram end of the queue. The remote system will also need to understand the encoding method and implement the necessary translations using the language of choice on the remote system.

The following pseudo code describes the struct being used in the operation. Fields `idNumber` and `dateOfBirth` will use the default conversion methods for their type and will be converted into ten and eight character strings respectively. The `historyText` field is a variable length field and will be encoded and decoded by means of a custom mapper class.

```
struct PersonHistory {
    String<10> idNumber;
    String historyText;
    Date dateOfBirth;
}
```

Figure 27. Pseudo code for the struct to be mapped:

Method `addToHistory` of class `LegacyBPO` sends a `PersonHistory` struct to a legacy system, the legacy system will append text to the variable length field `historyText` and return an updated copy of `PersonHistory`.

```
interface LegacyBPO {
    PersonHistory addToHistory(dtls PersonHistory);
}
```

Figure 28. Pseudo code for the BPO interface

Note that field `historyText` is being used in two cases - once in the parameter to operation `addToPersonHistory` and once in the return value from the operation. Therefore, the custom mapper class must be specified for each of these cases in `QueueConnectorFieldMappers.properties` (the lines are broken up for clarity).

```
LegacyBPO.addToPersonHistory.dtls.historyText=
    com.acme.mqutils.VariableStringMapper
LegacyBPO.addToPersonHistory.return.historyText=
    com.acme.mqutils.VariableStringMapper
```

Figure 29. The property file entries linking the fields to the mapper

The following listing shows the implementation of the custom mapper class.

```

package com.acme.mqutils;
// implementation for variable length string field mapper class
public class VariableStringMapper
extends MQFieldMapper {

    /**
     * The size of a prefix at the beginning of the string
     * which specifies the length of following data.
     */
    private static final int kStringHeaderInfoLength = 6;

    /**
     * Gets the encoded version of the mapped field within
     * the given struct.
     *
     * @param object The struct class containing the
     *             mapped field.
     * @return The field encoded as a String.
     * @throws AppException if the field could not be encoded.
     */
    public String encode(Object object)
    throws AppException {
        String historyText = null;
        // get the "historyText" field from the given struct:
        try {
            historyText = (String) getMappedField().get(object);
        } catch (IllegalAccessException e) {
            // use the handler in the superclass to deal with
            // this exception:
            handleEncodingException(e, object);
        }

        // construct the prefix which will hold the
        // size of the data.
        int bufferLength = historyText.length();

        String sizeSpecifierString = String.valueOf(bufferLength);
        // pad the size specifier to the right length
        sizeSpecifierString = MQUtils.padRight(
            sizeSpecifierString,
            kStringHeaderInfoLength);
        // put the prefix and the data together.
        String result = sizeSpecifierString + historyText;
        return result;
    }

    /**
     * Decodes the given string and assigns the resulting value
     * to the mapped field in the struct.
     *
     * @param object The struct class containing the field
     * @param encodedString The encoded form of the data.
     * @return the number of characters consumed from the
     *         encoded string.
     * @throws AppException if the target struct field could
     *         not be accessed.
     */
    public int decode(Object object, String encodedString)
    throws AppException {
        // the first N characters contain an expression
        // specifying the width of the encoded field.
        String sizeSpecifierString =
            encodedString.substring(0, kStringHeaderInfoLength);
        sizeSpecifierString = sizeSpecifierString.trim();
        int sizeOfString = Integer.valueOf(
            sizeSpecifierString).intValue();
        // Now that we know the size of the data, take that
        // many characters of data from the encoded string:
        String historyText = encodedString.substring(
            kStringHeaderInfoLength,
            kStringHeaderInfoLength + sizeOfString);
        // Update field "historyText" of the given struct:

```

Examples of MQSeries messages transmitted and received by this connector operation are:

- 10000361iw4 One.19700714
- 10000361iw9 One. Two.19700714
- 10000361iw16 One. Two. Three.19700714

Where the first 10 characters are the *idNumber* field, the last 8 characters are the *dateOfBirth* field and the middle section is the variable length *historyText* field, of which the first six characters specify the length of the data.

## Example 2 - Working with Lists

In the following example a custom field mapper class is used to implement encoding and decoding of a struct which aggregates a list of another struct. The list is encoded into a single string whereby the first 4 characters contain a number specifying the number of entries in the list and the remainder of the string consists of the encoded form of each struct as a fixed length string. The main purpose of this example is to illustrate how list aggregations are handled when implementing a custom mapper class.

As with the previous example, this example only shows the implementation at the Cúram end of the queue. The remote system will also need to understand the encoding method and implement the necessary translations using the language of choice on the remote system.

The following pseudo code describes the struct being used in the operation. Struct `PersonDtls` will be encoded as a fixed length 18 character string. Struct `PersonDtlsList` will be encoded by encoding each struct in its list, concatenating the results into a string, and prefixing the string with a six character string specifying the number of entries in the list.

```
struct PersonDtls {
    String<8> idNumber;
    String<10> surname;
}

struct PersonDtlsList {
    sequence <PersonDtls> dtls;
}
```

Figure 31. Pseudo code for the structs to be mapped:

Method `processNames` of class `LegacyBPO` sends a `PersonDtlsList` struct to a legacy system, the legacy system will perform some processing on this data and return an updated copy of `PersonDtlsList`.

```
interface LegacyBPO {
    PersonDtlsList processNames(p1 PersonDtlsList);
}
```

Figure 32. Pseudo code for the BPO interface

Again, as in the previous example, field `dtls` of struct `PersonDtlsList` is being used in two cases: once in the parameter to operation `processNames` and once in the return value from the operation. Therefore the custom mapper class must be specified for each of these cases in `QueueConnectorFieldMappers.properties` (the lines have been split for clarity).

```
LegacyBPO.processNames.p1.dtls=  
    com.acme.mqutils.PersonDtlsListMapper  
LegacyBPO.processNames.return.dtls=  
    com.acme.mqutils.PersonDtlsListMapper
```

*Figure 33. The property file entry linking the fields to the mapper*

The following listing shows the implementation of the custom mapper class.

```

package com.acme.mqutils;

// implementation
public class PersonDtIsListMapper {

    /**
     * The size of a prefix at the beginning of the string
     * which specifies the number of encoded entries in the
     * remainder of the string.
     */
    private static final int kStringHeaderInfoLength = 4;

    /**
     * The number of characters used to encode one
     * 'PersonDtIs' struct.
     */
    private static final int kLengthOfOneEncodedStruct = 18;

    /**
     * Encodes the 'dtIs' member into a string. The first 4
     * characters contain the number of items in the list, the
     * rest of the string consists of the encoded version of
     * each struct in the list concatenated together.
     *
     * @param object the object containing the field to be
     *             encoded
     * @throws AppException if it couldn't be encoded
     * @return A encoded string.
     */
    public String encode(Object object) throws AppException {

        PersonDtIsList.List_dtIs d = null;

        try {
            // get a reference to the field within the struct
            // to be encoded
            d = (PersonDtIsList.List_dtIs)
                getMappedField().get(object);
        } catch (IllegalAccessException e) {
            // use the handler in the superclass to deal with
            // this exception:
            handleEncodingException(e, object);
        }

        // construct the prefix which will specify the number
        // of items in the list.
        int bufferLength = d.size();
        String sizeSpecifierString =
            String.valueOf(bufferLength);

        // apply padding to make it the right size
        sizeSpecifierString =
            MQUtils.padRight(
                sizeSpecifierString, kStringHeaderInfoLength);

        // Now go through the items in the
        // list and encode each one.
        String data = "";
        for (int i = 0; i < d.size(); i++) {
            PersonDtIs currentItem = d.item(i);
            data += encodeOneEntry(currentItem);
        }

        // put the prefix and the data together.
        String result = sizeSpecifierString + data;
        return result;
    }
}

```

86 **\*\*** IBM Cúram Social Program Management; Cúram Modeling Reference Guide  
 \* Decodes a series of PersonDtIs entries in the string  
 \* and adds them to field PersonDtIsList.List\_dtIs in the  
 \* given PersonDtIsList object.  
 \*

For example, the following list of Ent18131 structs:

- ("0000361i", "James")
- ("0024684x", "John")
- ("8211519f", "Sharon")

would be encoded as follows:

```
"3 0000361iJames  
0024684xJohn 8211519fSharon "
```

where the first four characters contain a number specifying the number of encoded structs to follow, and the remaining string consists of three 18 character blocks corresponding to the three encoded structs.

---

## Subclassing

Use subclassing for process, facade, entity, webservice, and ws inbound classes. Subclassing allows you to add new functionality or override existing functionality.

Subclassing cannot be used to add extra attributes to entities or structs.

Reasons to subclass a class include:

- Adding new stereotyped methods to an existing entity class.
- Adding or contributing to an existing entity or operation exit points.
- Modifying existing entity operationReadmulti Max options.

## How to Model It

### Basic Subclassing

A class is transformed into a subclass by adding a “generalization” relationship from the subclass to the superclass (base class). On a class diagram this appears as a line between the two classes with an arrow pointing toward the superclass.

This means that the subclass inherits all the operations of the superclass, and in addition it may:

- Add extra functions;
- Modify the applicable options of the function in the superclass.

Consider two classes where MySubclass is a subclass of MyBaseClass:

- MyBaseClass has two operations: op1() and op2()
- MySubclass has three operations: op1(), op2() and op3() where op1 and op2 is inherited from MyBaseClass and MySubclass.op3 is provided only in the MySubclass class.

### Replacing the Superclass

When you define a subclass, you may specify that the subclass replaces its superclass entirely. To turn on the feature for an individual entity class the Replace\_Superclass property in the Rational Software Architect Curam Properties tab must be set to 1 - yes using the supplied drop-down.

For example, setting Replace\_Superclass to yes for a class, MySubclass, means that instances of the base class, MyBaseClass, will no longer be created. All requests for

the base class (MyBaseClass) will now receive an instance of the subclass (MySubClass). This is handled by the factory mechanism and is transparent to the user.

### **Abstract Classes**

A class is made abstract by setting its `Abstract` option to `yes` in the meta-model. In this case the generated Java class hierarchy for this abstract class will not include a factory class. This means that the class cannot be instantiated and the only purpose of having the class is to enable it to be subclassed.

All non-abstract subclasses of the abstract classes will have the factory component and are instantiated in the normal way.

The developer must provide the `impl` Java code for abstract classes (unless the abstract class has no subclasses). From here on the usual rules for abstract classes apply: the `impl` class can contain implementations for some or all of the methods declared in the class, and any methods for which an implementation has not been provided must be implemented by the subclass(es).

### **Restrictions**

- Multiple inheritance is not supported by the Cúram generators.
- Subclassing can only be used to add or override operations, it cannot be used to add or override attributes.

## **How to write Code for Subclassing**

There are no specific restrictions on writing code for subclassing. It is possible to subclass any entity, facade or process class without having to change the way that class is declared or used.

New subclasses of existing classes should be written in new source files. All new source files should be placed within the source subdirectory of the `EJBServer\components\<custom>` directory. Where `<custom>` is any new directory created under the `components` directory that conforms to the same directory structure as `components\core`. The generated class hierarchy will dictate the packaging of the new source files.

## **Example - Using Subclassing to Override Entity Exit Points**

### **Overriding Validation Exit Point**

In order to override the validation exit point of an entity in a subclass:

- enable the `Automatic validation operation` option on an entity subclass;
- specify at least one of the entity superclass stereotype `insert` or `modify` operations in the subclass.

For example, consider two classes, `MyEntityClass` and `MyEntitySubClass`. The subclass, `MyEntitySubClass`, would inherit the key and details of the superclass. `MyEntitySubClass` would have the `Automatic validation operation` option enabled and would add the `insert` or `modify` operations.

For more information on validation exit point, see “Validation” on page 21.

### **Overriding Pre Data Access, Post Data Access, and On-Fail exit points**

In order to override the `Pre Data Access`, `Post Data Access`, or `On-fail` exit points of an entity in a subclass:

- specify the operation(s) of entity superclass in the subclass;
- enable the Pre Data Access, Post Data Access, or On-fail options as appropriate on the operations of the entity subclass.

For example, consider two classes, `EntityClass` and `EntitySubClass` with the subclass, `EntitySubClass`, inheriting the key and details of the superclass. The same operations would be defined in both classes; e.g.: insert, read, and modify. In both these classes these operations would have the following exit point options enabled:

- On Fail operation is enabled on operation insert;
- Post Data Access operation is enabled on operation read;
- Pre Data Access operation is enabled on operation modify.

For more information on exit points see “Exit Points” on page 21.

---

## Application Customization

You should store customizations separately from the original model. You can then upgrade the original model without overwriting customizations.

One of the more difficult aspects of customizing an application is the handling of upgrades to the original model. Any changes which have been stored with the original model will be overwritten when a newer version of the model is taken on. This can be prevented by storing customizations separately from the original model to avoid overwriting customizations.

The following features are available to facilitate the customization of an application:

- “Extension Classes” on page 24;
- “Overriding a Domain Definition” on page 15;
- “Subclassing” on page 87.



---

## Notices

This information was developed for products and services offered in the United States.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-17*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihombashi*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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 and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

---

## Privacy Policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings

can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies or other similar technologies that collect each user's name, user name, password, and/or other personally identifiable information for purposes of session management, authentication, enhanced user usability, single sign-on configuration and/or other usage tracking and/or functional purposes. These cookies or other similar technologies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at " Copyright and trademark information " at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft 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.

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







Printed in USA