

Seam - Contextual Components

A Framework for Enterprise Java

Version: 2.0.1.CR1

Table of Contents

Introduction to JBoss Seam	xi
1. Seam Tutorial	1
1.1. Try the examples	1
1.1.1. Running the examples on JBoss AS	1
1.1.2. Running the examples on Tomcat	1
1.1.3. Running the example tests	1
1.2. Your first Seam application: the registration example	2
1.2.1. Understanding the code	2
1.2.1.1. The entity bean: User.java	3
1.2.1.2. The stateless session bean class: RegisterAction.java	5
1.2.1.3. The session bean local interface: Register.java	6
1.2.1.4. The Seam component deployment descriptor: components.xml	6
1.2.1.5. The web deployment description: web.xml	7
1.2.1.6. The JSF configuration: faces-config.xml	8
1.2.1.7. The EJB deployment descriptor: ejb-jar.xml	8
1.2.1.8. The EJB persistence deployment descriptor: persistence.xml	9
1.2.1.9. The view: register.xhtml and registered.xhtml	9
1.2.1.10. The EAR deployment descriptor: application.xml	10
1.2.2. How it works	11
1.3. Clickable lists in Seam: the messages example	11
1.3.1. Understanding the code	12
1.3.1.1. The entity bean: Message.java	12
1.3.1.2. The stateful session bean: MessageManagerBean.java	13
1.3.1.3. The session bean local interface: MessageManager.java	14
1.3.1.4. The view: messages.jsp	15
1.3.2. How it works	16
1.4. Seam and jBPM: the todo list example	16
1.4.1. Understanding the code	17
1.4.2. How it works	22
1.5. Seam pageflow: the numberguess example	22
1.5.1. Understanding the code	23
1.5.2. How it works	28
1.6. A complete Seam application: the Hotel Booking example	28
1.6.1. Introduction	28
1.6.2. Overview of the booking example	30
1.6.3. Understanding Seam conversations	30
1.6.4. The Seam UI control library	36
1.6.5. The Seam Debug Page	36
1.7. A complete application featuring Seam and jBPM: the DVD Store example	37
1.8. An example of Seam with Hibernate: the Hibernate Booking example	39
1.9. A RESTful Seam application: the Blog example	39
1.9.1. Using "pull"-style MVC	40
1.9.2. Bookmarkable search results page	41
1.9.3. Using "push"-style MVC in a RESTful application	44
2. Getting started with Seam, using seam-gen	47
2.1. Before you start	47
2.2. Setting up a new Eclipse project	47
2.3. Creating a new action	49

2.4. Creating a form with an action	50
2.5. Generating an application from an existing database	51
2.6. Generating an application from existing JPA/EJB3 entities	51
2.7. Deploying the application as an EAR	51
2.8. Seam and incremental hot deployment	52
2.9. Using Seam with JBoss 4.0	53
2.9.1. Install JBoss 4.0	53
2.9.2. Install the JSF 1.2 RI	53
3. The contextual component model	54
3.1. Seam contexts	54
3.1.1. Stateless context	54
3.1.2. Event context	54
3.1.3. Page context	55
3.1.4. Conversation context	55
3.1.5. Session context	55
3.1.6. Business process context	56
3.1.7. Application context	56
3.1.8. Context variables	56
3.1.9. Context search priority	56
3.1.10. Concurrency model	57
3.2. Seam components	57
3.2.1. Stateless session beans	58
3.2.2. Stateful session beans	58
3.2.3. Entity beans	58
3.2.4. JavaBeans	59
3.2.5. Message-driven beans	59
3.2.6. Interception	59
3.2.7. Component names	60
3.2.8. Defining the component scope	61
3.2.9. Components with multiple roles	61
3.2.10. Built-in components	61
3.3. Bijection	62
3.4. Lifecycle methods	64
3.5. Conditional installation	64
3.6. Logging	65
3.7. The Mutable interface and @ReadOnly	66
3.8. Factory and manager components	67
4. Configuring Seam components	70
4.1. Configuring components via property settings	70
4.2. Configuring components via components.xml	70
4.3. Fine-grained configuration files	73
4.4. Configurable property types	73
4.5. Using XML Namespaces	74
5. Events, interceptors and exception handling	77
5.1. Seam events	77
5.1.1. Page actions	77
5.1.1.1. Page parameters	78
5.1.1.2. Navigation	80
5.1.1.3. Fine-grained files for definition of navigation, page actions and parameters	82
5.1.2. Component-driven events	82
5.1.3. Contextual events	84
5.2. Seam interceptors	85

5.3. Managing exceptions	86
5.3.1. Exceptions and transactions	87
5.3.2. Enabling Seam exception handling	87
5.3.3. Using annotations for exception handling	87
5.3.4. Using XML for exception handling	88
5.3.5. Some common exceptions	89
6. Conversations and workspace management	90
6.1. Seam's conversation model	90
6.2. Nested conversations	92
6.3. Starting conversations with GET requests	92
6.4. Using <s:link> and <s:button>	93
6.5. Success messages	94
6.6. Using an "explicit" conversation id	95
6.7. Natural conversation ids	96
6.7.1. Configuration	96
6.7.2. Redirecting to a natural conversation	96
6.8. Workspace management	97
6.8.1. Workspace management and JSF navigation	97
6.8.2. Workspace management and jPDL pageflow	98
6.8.3. The conversation switcher	98
6.8.4. The conversation list	99
6.8.5. Breadcrumbs	99
6.9. Conversational components and JSF component bindings	99
6.10. Concurrent calls to conversational components	100
6.10.1. RichFaces Ajax	101
7. Pageflows and business processes	103
7.1. Pageflow in Seam	103
7.1.1. The two navigation models	103
7.1.2. Seam and the back button	106
7.2. Using jPDL pageflows	107
7.2.1. Installing pageflows	107
7.2.2. Starting pageflows	107
7.2.3. Page nodes and transitions	108
7.2.4. Controlling the flow	109
7.2.5. Ending the flow	109
7.2.6. Pageflow composition	109
7.3. Business process management in Seam	110
7.4. Using jPDL business process definitions	110
7.4.1. Installing process definitions	111
7.4.2. Initializing actor ids	111
7.4.3. Initiating a business process	111
7.4.4. Task assignment	111
7.4.5. Task lists	111
7.4.6. Performing a task	112
8. Seam and Object/Relational Mapping	114
8.1. Introduction	114
8.2. Seam managed transactions	114
8.2.1. Disabling Seam-managed transactions	115
8.2.2. Configuring a Seam transaction manager	115
8.2.3. Transaction synchronization	116
8.3. Seam-managed persistence contexts	116
8.3.1. Using a Seam-managed persistence context with JPA	117

8.3.2. Using a Seam-managed Hibernate session	117
8.3.3. Seam-managed persistence contexts and atomic conversations	118
8.4. Using the JPA "delegate"	119
8.5. Using EL in EJB-QL/HQL	119
8.6. Using Hibernate filters	120
9. JSF form validation in Seam	121
10. Groovy integration	125
10.1. Groovy introduction	125
10.2. Writing Seam applications in Groovy	125
10.2.1. Writing Groovy components	125
10.2.1.1. Entity	125
10.2.1.2. Seam component	126
10.2.2. seam-gen	127
10.3. Deployment	127
10.3.1. Deploying Groovy code	127
10.3.2. Native .groovy file deployment at development time	127
10.3.3. seam-gen	127
11. The Seam Application Framework	129
11.1. Introduction	129
11.2. Home objects	130
11.3. Query objects	133
11.4. Controller objects	135
12. Seam and JBoss Rules	136
12.1. Installing rules	136
12.2. Using rules from a Seam component	137
12.3. Using rules from a jBPM process definition	137
13. Security	139
13.1. Overview	139
13.1.1. Which mode is right for my application?	139
13.2. Requirements	139
13.3. Disabling Security	139
13.4. Authentication	140
13.4.1. Configuration	140
13.4.2. Writing an authentication method	140
13.4.2.1. Identity.addRole()	141
13.4.2.2. Special Considerations	142
13.4.3. Writing a login form	142
13.4.4. Simplified Configuration - Summary	143
13.4.5. Handling Security Exceptions	143
13.4.6. Login Redirection	144
13.4.7. HTTP Authentication	144
13.4.7.1. Writing a Digest Authenticator	145
13.4.8. Advanced Authentication Features	145
13.4.8.1. Using your container's JAAS configuration	145
13.5. Error Messages	145
13.6. Authorization	146
13.6.1. Core concepts	146
13.6.2. Securing components	146
13.6.2.1. The @Restrict annotation	146
13.6.2.2. Inline restrictions	147
13.6.3. Security in the user interface	148
13.6.4. Securing pages	148

13.6.5. Securing Entities	149
13.6.5.1. Entity security with JPA	150
13.6.5.2. Entity security with Hibernate	150
13.7. Writing Security Rules	150
13.7.1. Permissions Overview	151
13.7.2. Configuring a rules file	151
13.7.3. Creating a security rules file	151
13.7.3.1. Wildcard permission checks	153
13.8. SSL Security	153
13.9. CAPTCHA	154
13.9.1. Configuring the CAPTCHA Servlet	154
13.9.2. Adding a CAPTCHA to a form	154
13.9.3. Customising the CAPTCHA algorithm	154
13.10. Security Events	155
13.11. Extending the Identity component	155
14. Internationalization and themes	157
14.1. Locales	157
14.2. Labels	157
14.2.1. Defining labels	158
14.2.2. Displaying labels	158
14.2.3. Faces messages	159
14.3. Timezones	159
14.4. Themes	159
14.5. Persisting locale and theme preferences via cookies	160
15. Seam Text	161
15.1. Basic fomatting	161
15.2. Entering code and text with special characters	162
15.3. Links	163
15.4. Entering HTML	163
16. iText PDF generation	164
16.1. Using PDF Support	164
16.1.1. Creating a document	164
16.1.2. Basic Text Elements	165
16.1.3. Headers and Footers	168
16.1.4. Chapters and Sections	169
16.1.5. Lists	170
16.1.6. Tables	171
16.1.7. Document Constants	173
16.1.7.1. Color Values	174
16.1.7.2. Alignment Values	174
16.1.8. Configuring iText	174
16.2. Charting	175
16.3. Bar codes	181
16.4. Further documentation	182
17. Email	183
17.1. Creating a message	183
17.1.1. Attachments	184
17.1.2. HTML/Text alternative part	185
17.1.3. Multiple recipients	185
17.1.4. Multiple messages	185
17.1.5. Templating	185
17.1.6. Internationalisation	186

17.1.7. Other Headers	186
17.2. Receiving emails	186
17.3. Configuration	187
17.3.1. mailSession	187
17.3.1.1. JNDI lookup in JBoss AS	187
17.3.1.2. Seam configured Session	188
17.4. Meldware	188
17.5. Tags	188
18. Asynchronicity and messaging	191
18.1. Asynchronicity	191
18.1.1. Asynchronous methods	191
18.1.2. Asynchronous methods with the Quartz Dispatcher	193
18.1.3. Asynchronous events	195
18.2. Messaging in Seam	195
18.2.1. Configuration	195
18.2.2. Sending messages	196
18.2.3. Receiving messages using a message-driven bean	196
18.2.4. Receiving messages in the client	196
19. Caching	198
19.1. Using JBossCache in Seam	199
19.2. Page fragment caching	199
20. Web Services	201
20.1. Configuration and Packaging	201
20.2. Conversational Web Services	201
20.2.1. A Recommended Strategy	202
20.3. An example web service	202
21. Remoting	204
21.1. Configuration	204
21.2. The "Seam" object	204
21.2.1. A Hello World example	205
21.2.2. Seam.Component	206
21.2.2.1. Seam.Component.newInstance()	206
21.2.2.2. Seam.Component.getInstance()	207
21.2.2.3. Seam.Component.getComponentName()	207
21.2.3. Seam.Remoting	207
21.2.3.1. Seam.Remoting.createType()	207
21.2.3.2. Seam.Remoting.getTypeName()	207
21.3. Evaluating EL Expressions	208
21.4. Client Interfaces	208
21.5. The Context	208
21.5.1. Setting and reading the Conversation ID	209
21.5.2. Remote calls within the current conversation scope	209
21.6. Batch Requests	209
21.7. Working with Data types	209
21.7.1. Primitives / Basic Types	209
21.7.1.1. String	210
21.7.1.2. Number	210
21.7.1.3. Boolean	210
21.7.2. JavaBeans	210
21.7.3. Dates and Times	210
21.7.4. Enums	210
21.7.5. Collections	211

21.7.5.1. Bags	211
21.7.5.2. Maps	211
21.8. Debugging	211
21.9. The Loading Message	212
21.9.1. Changing the message	212
21.9.2. Hiding the loading message	212
21.9.3. A Custom Loading Indicator	212
21.10. Controlling what data is returned	212
21.10.1. Constraining normal fields	213
21.10.2. Constraining Maps and Collections	213
21.10.3. Constraining objects of a specific type	213
21.10.4. Combining Constraints	214
21.11. JMS Messaging	214
21.11.1. Configuration	214
21.11.2. Subscribing to a JMS Topic	214
21.11.3. Unsubscribing from a Topic	214
21.11.4. Tuning the Polling Process	214
22. Seam and the Google Web Toolkit	216
22.1. Configuration	216
22.2. Preparing your component	216
22.3. Hooking up a GWT widget to the Seam component	217
22.4. GWT Ant Targets	218
23. Spring Framework integration	220
23.1. Injecting Seam components into Spring beans	220
23.2. Injecting Spring beans into Seam components	221
23.3. Making a Spring bean into a Seam component	222
23.4. Seam-scoped Spring beans	222
23.5. Using Spring PlatformTransactionManagement	223
23.6. Using a Seam Managed Persistence Context in Spring	223
23.7. Using a Seam Managed Hibernate Session in Spring	224
23.8. Spring Application Context as a Seam Component	225
23.9. Using a Spring TaskExecutor for @Asynchronous	225
24. Hibernate Search	227
24.1. Introduction	227
24.2. Configuration	227
24.3. Usage	228
25. Configuring Seam and packaging Seam applications	230
25.1. Basic Seam configuration	230
25.1.1. Integrating Seam with JSF and your servlet container	230
25.1.2. Using facelets	230
25.1.3. Seam Resource Servlet	231
25.1.4. Seam servlet filters	231
25.1.4.1. Exception handling	231
25.1.4.2. Conversation propagation with redirects	232
25.1.4.3. Multipart form submissions	232
25.1.4.4. Character encoding	232
25.1.4.5. RichFaces	233
25.1.4.6. Identity Logging	233
25.1.4.7. Context management for custom servlets	233
25.1.4.8. Adding custom filters	234
25.1.5. Integrating Seam with your EJB container	234
25.1.6. Don't forget!	235

25.2. Configuring Seam in Java EE 5	235
25.2.1. Packaging	236
25.3. Configuring Seam in J2EE	237
25.3.1. Bootstrapping Hibernate in Seam	237
25.3.2. Bootstrapping JPA in Seam	237
25.3.3. Packaging	238
25.4. Configuring Seam in Java SE, without JBoss Embedded	238
25.5. Configuring Seam in Java SE, with JBoss Embedded	239
25.5.1. Installing Embedded JBoss	239
25.5.2. Packaging	240
25.6. Configuring jBPM in Seam	240
25.6.1. Packaging	241
25.7. Configuring Seam in a Portal	242
25.8. Configuring SFSB and Session Timeouts in JBoss AS	242
26. Seam on OC4J	244
26.1. The jee5/booking example	244
26.1.1. Booking Example Dependencies	244
26.1.2. Extra dependencies required by OC4J	244
26.1.3. Configuration file changes	245
26.1.4. Building the jee5/booking example	245
26.2. Deploying a Seam application to OC4J	246
26.3. Deploying an application created using seam-gen to OC4J	247
26.3.1. OC4J Deployment Descriptors for the seam-gen'd application	251
27. Seam annotations	252
27.1. Annotations for component definition	252
27.2. Annotations for bijection	254
27.3. Annotations for component lifecycle methods	257
27.4. Annotations for context demarcation	257
27.5. Annotations for use with Seam JavaBean components in a J2EE environment	260
27.6. Annotations for exceptions	261
27.7. Annotations for Seam Remoting	261
27.8. Annotations for Seam interceptors	262
27.9. Annotations for asynchronicity	262
27.10. Annotations for use with JSF	263
27.10.1. Annotations for use with dataTable	263
27.11. Meta-annotations for databinding	264
27.12. Annotations for packaging	265
27.13. Annotations for integrating with the servlet container	265
28. Built-in Seam components	266
28.1. Context injection components	266
28.2. Utility components	266
28.3. Components for internationalization and themes	268
28.4. Components for controlling conversations	269
28.5. jBPM-related components	270
28.6. Security-related components	272
28.7. JMS-related components	272
28.8. Mail-related components	272
28.9. Infrastructural components	272
28.10. Miscellaneous components	275
28.11. Special components	275
29. Seam JSF controls	277
29.1. Tags	277

29.2. Annotations	290
30. Expression language enhancements	292
30.1. Parameterized Method Bindings	292
30.1.1. Usage	292
30.1.2. Limitations	292
30.1.2.1. Incompatibility with JSP 2.1	293
30.1.2.2. Calling a MethodExpression from Java code	293
30.2. Parameterized Value Bindings	293
30.3. Projection	293
31. Testing Seam applications	295
31.1. Unit testing Seam components	295
31.2. Integration testing Seam components	296
31.2.1. Using mocks in integration tests	297
31.3. Integration testing Seam application user interactions	297
31.3.1. Integration Testing with Mock Data	300
31.3.2. Integration Testing Seam Mail	301
32. Seam tools	303
32.1. jBPM designer and viewer	303
32.1.1. Business process designer	303
32.1.2. Pageflow viewer	303
33. Dependencies	305
33.1. Project Dependencies	305
33.1.1. Core	305
33.1.2. RichFaces	306
33.1.3. Seam Mail	306
33.1.4. Seam PDF	307
33.1.5. JBoss Rules	307
33.1.6. JBPM	307
33.1.7. GWT	308
33.1.8. Spring	308
33.1.9. Groovy	308
33.2. Dependency Management using Maven	308

Introduction to JBoss Seam

Seam is an application framework for Enterprise Java. It is inspired by the following principles:

One kind of "stuff"

Seam defines a uniform component model for all business logic in your application. A Seam component may be stateful, with the state associated with any one of several well-defined contexts, including the long-running, persistent, *business process context* and the *conversation context*, which is preserved across multiple web requests in a user interaction.

There is no distinction between presentation tier components and business logic components in Seam. You can layer your application according to whatever architecture you devise, rather than being forced to shoe-horn your application logic into an unnatural layering scheme forced upon you by whatever combination of stovepipe frameworks you're using today.

Unlike plain Java EE or J2EE components, Seam components may *simultaneously* access state associated with the web request and state held in transactional resources (without the need to propagate web request state manually via method parameters). You might object that the application layering imposed upon you by the old J2EE platform was a Good Thing. Well, nothing stops you creating an equivalent layered architecture using Seam—the difference is that *you* get to architect your own application and decide what the layers are and how they work together.

Integrate JSF with EJB 3.0

JSF and EJB 3.0 are two of the best new features of Java EE 5. EJB3 is a brand new component model for server side business and persistence logic. Meanwhile, JSF is a great component model for the presentation tier. Unfortunately, neither component model is able to solve all problems in computing by itself. Indeed, JSF and EJB3 work best used together. But the Java EE 5 specification provides no standard way to integrate the two component models. Fortunately, the creators of both models foresaw this situation and provided standard extension points to allow extension and integration with other frameworks.

Seam unifies the component models of JSF and EJB3, eliminating glue code, and letting the developer think about the business problem.

It is possible to write Seam applications where "everything" is an EJB. This may come as a surprise if you're used to thinking of EJBs as coarse-grained, so-called "heavyweight" objects. However, version 3.0 has completely changed the nature of EJB from the point of view of the developer. An EJB is a fine-grained object—nothing more complex than an annotated JavaBean. Seam even encourages you to use session beans as JSF action listeners!

On the other hand, if you prefer not to adopt EJB 3.0 at this time, you don't have to. Virtually any Java class may be a Seam component, and Seam provides all the functionality that you expect from a "light-weight" container, and more, for any component, EJB or otherwise.

Integrated AJAX

Seam supports the best open source JSF-based AJAX solutions: JBoss RichFaces and ICEfaces. These solutions let you add AJAX capability to your user interface without the need to write any JavaScript code.

Alternatively, Seam provides a built-in JavaScript remoting layer that lets you call components asynchronously from client-side JavaScript without the need for an intermediate action layer. You can even subscribe to server-side JMS topics and receive messages via AJAX push.

Neither of these approaches would work well, were it not for Seam's built-in concurrency and state man-

agement, which ensures that many concurrent fine-grained, asynchronous AJAX requests are handled safely and efficiently on the server side.

Business process as a first class construct

Optionally, Seam provides transparent business process management via jBPM. You won't believe how easy it is to implement complex workflows, collaboration and task management using jBPM and Seam.

Seam even allows you to define presentation tier pageflow using the same language (jPDL) that jBPM uses for business process definition.

JSF provides an incredibly rich event model for the presentation tier. Seam enhances this model by exposing jBPM's business process related events via exactly the same event handling mechanism, providing a uniform event model for Seam's uniform component model.

Declarative state management

We're all used to the concept of declarative transaction management and declarative security from the early days of EJB. EJB 3.0 even introduces declarative persistence context management. These are three examples of a broader problem of managing state that is associated with a particular *context*, while ensuring that all needed cleanup occurs when the context ends. Seam takes the concept of declarative state management much further and applies it to *application state*. Traditionally, J2EE applications implement state management manually, by getting and setting servlet session and request attributes. This approach to state management is the source of many bugs and memory leaks when applications fail to clean up session attributes, or when session data associated with different workflows collides in a multi-window application. Seam has the potential to almost entirely eliminate this class of bugs.

Declarative application state management is made possible by the richness of the *context model* defined by Seam. Seam extends the context model defined by the servlet spec—request, session, application—with two new contexts—conversation and business process—that are more meaningful from the point of view of the business logic.

You'll be amazed at how many things become easier once you start using conversations. Have you ever suffered pain dealing with lazy association fetching in an ORM solution like Hibernate or JPA? Seam's conversation-scoped persistence contexts mean you'll rarely have to see a `LazyInitializationException`. Have you ever had problems with the refresh button? The back button? With duplicate form submission? With propagating messages across a post-then-redirect? Seam's conversation management solves these problems without you even needing to really think about them. They're all symptoms of the broken state management architecture has been prevalent since the earliest days of the web.

Bijection

The notion of *Inversion of Control* or *dependency injection* exists in both JSF and EJB3, as well as in numerous so-called "lightweight containers". Most of these containers emphasize injection of components that implement *stateless services*. Even when injection of stateful components is supported (such as in JSF), it is virtually useless for handling application state because the scope of the stateful component cannot be defined with sufficient flexibility, and because components belonging to wider scopes may not be injected into components belonging to narrower scopes.

Bijection differs from IoC in that it is *dynamic*, *contextual*, and *bidirectional*. You can think of it as a mechanism for aliasing contextual variables (names in the various contexts bound to the current thread) to attributes of the component. Bijection allows auto-assembly of stateful components by the container. It even allows a component to safely and easily manipulate the value of a context variable, just by assigning to an attribute of the component.

Workspace management and multi-window browsing

Seam applications let the user freely switch between multiple browser tabs, each associated with a differ-

ent, safely isolated, conversation. Applications may even take advantage of *workspace management*, allowing the user to switch between conversations (workspaces) in a single browser tab. Seam provides not only correct multi-window operation, but also multi-window-like operation in a single window!

Prefer annotations to XML

Traditionally, the Java community has been in a state of deep confusion about precisely what kinds of meta-information counts as configuration. J2EE and popular "lightweight" containers have provided XML-based deployment descriptors both for things which are truly configurable between different deployments of the system, and for any other kinds of declaration which can not easily be expressed in Java. Java 5 annotations changed all this.

EJB 3.0 embraces annotations and "configuration by exception" as the easiest way to provide information to the container in a declarative form. Unfortunately, JSF is still heavily dependent on verbose XML configuration files. Seam extends the annotations provided by EJB 3.0 with a set of annotations for declarative state management and declarative context demarcation. This lets you eliminate the noisy JSF managed bean declarations and reduce the required XML to just that information which truly belongs in XML (the JSF navigation rules).

Integration testing is easy

Seam components, being plain Java classes, are by nature unit testable. But for complex applications, unit testing alone is insufficient. Integration testing has traditionally been a messy and difficult task for Java web applications. Therefore, Seam provides for testability of Seam applications as a core feature of the framework. You can easily write JUnit or TestNG tests that reproduce a whole interaction with a user, exercising all components of the system apart from the view (the JSP or Facelets page). You can run these tests directly inside your IDE, where Seam will automatically deploy EJB components using JBoss Embedded.

The specs ain't perfect

We think the latest incarnation of Java EE is great. But we know it's never going to be perfect. Where there are holes in the specifications (for example, limitations in the JSF lifecycle for GET requests), Seam fixes them. And the authors of Seam are working with the JCP expert groups to make sure those fixes make their way back into the next revision of the standards.


There's more to a web application than serving HTML pages

Today's web frameworks think too small. They let you get user input off a form and into your Java objects. And then they leave you hanging. A truly complete web application framework should address problems like persistence, concurrency, asynchronicity, state management, security, email, messaging, PDF and chart generation, workflow, wikitext rendering, webservices, caching and more. Once you scratch the surface of Seam, you'll be amazed at how many problems become simpler...

Seam integrates JPA and Hibernate3 for persistence, the EJB Timer Service and Quartz for lightweight asynchronicity, jBPM for workflow, JBoss Rules for business rules, Meldware Mail for email, Hibernate Search and Lucene for full text search, JMS for messaging and JBoss Cache for page fragment caching. Seam layers an innovative rule-based security framework over JAAS and JBoss Rules. There's even JSF tag libraries for rendering PDF, outgoing email, charts and wikitext. Seam components may be called synchronously as a Web Service, asynchronously from client-side JavaScript or Google Web Toolkit or, of course, directly from JSF.

Get started now!

Seam works in any Java EE application server, and even works in Tomcat. If your environment supports EJB 3.0, great! If it doesn't, no problem, you can use Seam's built-in transaction management with JPA or Hibernate3 for persistence. Or, you can deploy JBoss Embedded in Tomcat, and get full support for EJB 3.0.



It turns out that the combination of Seam, JSF and EJB3 is *the* simplest way to write a complex web application in Java. You won't believe how little code is required!

Chapter 1. Seam Tutorial

1.1. Try the examples

In this tutorial, we'll assume that you have downloaded JBoss AS 4.2.0. You should also have a copy of Seam downloaded and extracted to a work directory.

The directory structure of each example in Seam follows this pattern:

- Web pages, images and stylesheets may be found in `examples/registration/view`
- Resources such as deployment descriptors and data import scripts may be found in `examples/registration/resources`
- Java source code may be found in `examples/registration/src`
- The Ant build script is `examples/registration/build.xml`

1.1.1. Running the examples on JBoss AS

First, make sure you have Ant correctly installed, with `$ANT_HOME` and `$JAVA_HOME` set correctly. Next, make sure you set the location of your JBoss AS 4.2.0 installation in the `build.properties` file in the root folder of your Seam installation. If you haven't already done so, start JBoss AS now by typing `bin/run.sh` or `bin/run.bat` in the root directory of your JBoss installation.

Now, build and deploy the example by typing `ant deploy` in the `examples/registration` directory.

Try it out by accessing `http://localhost:8080/seam-registration/` [`http://localhost:8080/seam-registration/`] with your web browser.

1.1.2. Running the examples on Tomcat

First, make sure you have Ant correctly installed, with `$ANT_HOME` and `$JAVA_HOME` set correctly. Next, make sure you set the location of your Tomcat 6.0 installation in the `build.properties` file in the root folder of your Seam installation. You will need to follow the instructions in Section 25.5.1, “Installing Embedded JBoss” for installing JBoss Embedded on Tomcat 6.0. JBoss Embedded is required to run the Seam demo applications on Tomcat. (However, it is possible to use Seam on Tomcat without JBoss Embedded.)

Now, build and deploy the example by typing `ant tomcat.deploy` in the `examples/registration` directory.

Finally, start Tomcat.

Try it out by accessing `http://localhost:8080/jboss-seam-registration/` [`http://localhost:8080/jboss-seam-registration/`] with your web browser.

When you deploy the example to Tomcat, any EJB3 components will run inside the JBoss Embeddable EJB3 container, a complete standalone EJB3 container environment.

1.1.3. Running the example tests


Most of the examples come with a suite of TestNG integration tests. The easiest way to run the tests is to run `ant testexample` inside the `examples/registration` directory. It is also possible to run the tests inside your IDE using the TestNG plugin.

1.2. Your first Seam application: the registration example

The registration example is a fairly trivial application that lets a new user store his username, real name and password in the database. The example isn't intended to show off all of the cool functionality of Seam. However, it demonstrates the use of an EJB3 session bean as a JSF action listener, and basic configuration of Seam.

We'll go slowly, since we realize you might not yet be familiar with EJB 3.0.

The start page displays a very basic form with three input fields. Try filling them in and then submitting the form. This will save a user object in the database.




The screenshot shows a Mozilla Firefox browser window titled "Register New User - Mozilla Firefox". The address bar displays the URL "http://localhost:8080/seam-registration/register.seam". The browser's tab bar shows three tabs: "Chapter 1. Seam Tutorial", "Register New User" (which is the active tab), and "JBoss DVD Store". The main content area of the browser displays a registration form with the following fields and values:

- Username: gavin
- Real Name: Gavin King
- Password: *****

Below the password field is a "Register" button. The status bar at the bottom of the browser window indicates "Done".

1.2.1. Understanding the code

The example is implemented with two JSP pages, one entity bean and one stateless session bean.



Let's take a look at the code, starting from the "bottom".

1.2.1.1. The entity bean: `User.java`

We need an EJB entity bean for user data. This class defines *persistence* and *validation* declaratively, via annotations. It also needs some extra annotations that define the class as a Seam component.

Example 1.1.

```

@Entity (1)
@Name("user") (2)
@Scope(SESSION) (3)
@Table(name="users") (4)
public class User implements Serializable
{
    private static final long serialVersionUID = 1881413500711441951L;

    private String username; (5)
    private String password;
    private String name;

    public User(String name, String password, String username)
    {
        this.name = name;
        this.password = password;
        this.username = username;
    }

    public User() {} (6)

    @NotNull @Length(min=5, max=15) (7)
    public String getPassword()

```

```

    {
        return password;
    }

    public void setPassword(String password)
    {
        this.password = password;
    }

    @NotNull
    public String getName()
    {
        return name;
    }

    public void setName(String name)
    {
        this.name = name;
    }

    @Id @NotNull @Length(min=5, max=15)
    public String getUsername()
    {
        return username;
    }

    public void setUsername(String username)
    {
        this.username = username;
    }
}

```

(8)

- (1) The EJB3 standard `@Entity` annotation indicates that the `User` class is an entity bean.
- (2) A Seam component needs a *component name* specified by the `@Name` annotation. This name must be unique within the Seam application. When JSF asks Seam to resolve a context variable with a name that is the same as a Seam component name, and the context variable is currently undefined (null), Seam will instantiate that component, and bind the new instance to the context variable. In this case, Seam will instantiate a `User` the first time JSF encounters a variable named `user`.
- (3) Whenever Seam instantiates a component, it binds the new instance to a context variable in the component's *default context*. The default context is specified using the `@Scope` annotation. The `User` bean is a session scoped component.
- (4) The EJB standard `@Table` annotation indicates that the `User` class is mapped to the `users` table.
- (5) `name`, `password` and `username` are the persistent attributes of the entity bean. All of our persistent attributes define accessor methods. These are needed when this component is used by JSF in the render response and update model values phases.
- (6) An empty constructor is both required by both the EJB specification and by Seam.
- (7) The `@NotNull` and `@Length` annotations are part of the Hibernate Validator framework. Seam integrates Hibernate Validator and lets you use it for data validation (even if you are not using Hibernate for persistence).
- (8) The EJB standard `@Id` annotation indicates the primary key attribute of the entity bean.

The most important things to notice in this example are the `@Name` and `@Scope` annotations. These annotations establish that this class is a Seam component.

We'll see below that the properties of our `User` class are bound directly to JSF components and are populated by JSF during the update model values phase. We don't need any tedious glue code to copy data back and forth between the JSP pages and the entity bean domain model.

However, entity beans shouldn't do transaction management or database access. So we can't use this component as a JSF action listener. For that we need a session bean.

1.2.1.2. The stateless session bean class: RegisterAction.java

Most Seam application use session beans as JSF action listeners (you can use JavaBeans instead if you like).

We have exactly one JSF action in our application, and one session bean method attached to it. In this case, we'll use a stateless session bean, since all the state associated with our action is held by the User bean.

This is the only really interesting code in the example!

Example 1.2.

```

@Stateless                                     (1)
@Name("register")
public class RegisterAction implements Register
{
    @In                                         (2)
    private User user;

    @PersistenceContext                         (3)
    private EntityManager em;

    @Logger                                    (4)
    private Log log;

    public String register()                   (5)
    {
        List existing = em.createQuery(
            "select username from User where username=#{user.username}") (6)
            .getResultList();

        if (existing.size()==0)
        {
            em.persist(user);
            log.info("Registered new user #{user.username}");           (7)
            return "/registered.xhtml";                                   (8)
        }
        else
        {
            FacesMessages.instance().add("User #{user.username} already exists"); (9)
            return null;
        }
    }
}

```

- (1) The EJB standard `@Stateless` annotation marks this class as stateless session bean.
- (2) The `@In` annotation marks an attribute of the bean as injected by Seam. In this case, the attribute is injected from a context variable named `user` (the instance variable name).
- (3) The EJB standard `@PersistenceContext` annotation is used to inject the EJB3 entity manager.
- (4) The Seam `@Logger` annotation is used to inject the component's `Log` instance.
- (5) The action listener method uses the standard EJB3 `EntityManager` API to interact with the database, and returns the JSF outcome. Note that, since this is a session bean, a transaction is automatically begun when the `register()` method is called, and committed when it completes.
- (6) Notice that Seam lets you use a JSF EL expression inside EJB-QL. Under the covers, this results in an ordinary JPA `setParameter()` call on the standard JPA Query object. Nice, huh?
- (7) The `Log` API lets us easily display templated log messages.
- (8) JSF action listener methods return a string-valued outcome that determines what page will be displayed next. A null outcome (or a void action listener method) redisplay the previous page. In plain JSF, it is normal to always use a JSF *navigation rule* to determine the JSF view id from the outcome. For complex

application this indirection is useful and a good practice. However, for very simple examples like this one, Seam lets you use the JSF view id as the outcome, eliminating the requirement for a navigation rule. *Note that when you use a view id as an outcome, Seam always performs a browser redirect.*

- (9) Seam provides a number of *built-in components* to help solve common problems. The `FacesMessages` component makes it easy to display templated error or success messages. Built-in Seam components may be obtained by injection, or by calling an `instance()` method.

Note that we did not explicitly specify a `@Scope` this time. Each Seam component type has a default scope if not explicitly specified. For stateless session beans, the default scope is the stateless context. Actually, *all* stateless session beans belong in the stateless context.

Our session bean action listener performs the business and persistence logic for our mini-application. In more complex applications, we might need to layer the code and refactor persistence logic into a dedicated data access component. That's perfectly trivial to do. But notice that Seam does not force you into any particular strategy for application layering.

Furthermore, notice that our session bean has simultaneous access to context associated with the web request (the form values in the `User` object, for example), and state held in transactional resources (the `EntityManager` object). This is a break from traditional J2EE architectures. Again, if you are more comfortable with the traditional J2EE layering, you can certainly implement that in a Seam application. But for many applications, it's simply not very useful.

1.2.1.3. The session bean local interface: `Register.java`

Naturally, our session bean needs a local interface.

Example 1.3.

```
@Local
public interface Register
{
    public String register();
}
```

That's the end of the Java code. Now onto the deployment descriptors.

1.2.1.4. The Seam component deployment descriptor: `components.xml`

If you've used many Java frameworks before, you'll be used to having to declare all your component classes in some kind of XML file that gradually grows more and more unmanageable as your project matures. You'll be relieved to know that Seam does not require that application components be accompanied by XML. Most Seam applications require a very small amount of XML that does not grow very much as the project gets bigger.

Nevertheless, it is often useful to be able to provide for *some* external configuration of *some* components (particularly the components built in to Seam). You have a couple of options here, but the most flexible option is to provide this configuration in a file called `components.xml`, located in the `WEB-INF` directory. We'll use the `components.xml` file to tell Seam how to find our EJB components in JNDI:

Example 1.4.

```
<?xml version="1.0" encoding="UTF-8"?>
<components xmlns="http://jboss.com/products/seam/components"
```

```

xmlns:core="http://jboss.com/products/seam/core"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation=
    "http://jboss.com/products/seam/core http://jboss.com/products/seam/core-2.0.xsd
    http://jboss.com/products/seam/components http://jboss.com/products/seam/components-2.0.xsd"

    <core:init jndi-pattern="@jndiPattern@"/>

</components>

```

This code configures a property named `jndiPattern` of a built-in Seam component named `org.jboss.seam.core.init`. The funny `@` symbols are there because our Ant build script puts the correct JNDI pattern in when we deploy the application.

1.2.1.5. The web deployment description: `web.xml`

The presentation layer for our mini-application will be deployed in a WAR. So we'll need a web deployment descriptor.

Example 1.5.

```

<?xml version="1.0" encoding="UTF-8"?>
<web-app version="2.5"
  xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
    http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd">

  <!-- Seam -->

  <listener>
    <listener-class>org.jboss.seam.servlet.SeamListener</listener-class>
  </listener>

  <!-- JSF -->

  <listener>
    <listener-class>com.sun.faces.config.ConfigureListener</listener-class>
  </listener>

  <context-param>
    <param-name>javax.faces.DEFAULT_SUFFIX</param-name>
    <param-value>.xhtml</param-value>
  </context-param>

  <servlet>
    <servlet-name>Faces Servlet</servlet-name>
    <servlet-class>javax.faces.webapp.FacesServlet</servlet-class>
    <load-on-startup>1</load-on-startup>
  </servlet>

  <servlet-mapping>
    <servlet-name>Faces Servlet</servlet-name>
    <url-pattern>*.seam</url-pattern>
  </servlet-mapping>

  <session-config>
    <session-timeout>10</session-timeout>
  </session-config>

</web-app>

```

This `web.xml` file configures Seam and JSF. The configuration you see here is pretty much identical in all Seam applications.

1.2.1.6. The JSF configuration: `faces-config.xml`

Most Seam applications use JSF views as the presentation layer. So usually we'll need `faces-config.xml`. In our case, we are going to use Facelets for defining our views, so we need to tell JSF to use Facelets as its templating engine.

Example 1.6.

```
<?xml version="1.0" encoding="UTF-8"?>
<faces-config version="1.2"
    xmlns="http://java.sun.com/xml/ns/javaee"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://java.sun.com/xml/ns/javaee http://java.sun.com/xml/ns/javaee/
    http://java.sun.com/xml/ns/javaee/2006-03-08/jsf-config_1_2.xsd"
    >

    <!-- Facelets support -->
    <application>
        <view-handler>com.sun.facelets.FaceletViewHandler</view-handler>
    </application>

</faces-config>
```

Note that we don't need any JSF managed bean declarations! Our managed beans are annotated Seam components. In Seam applications, the `faces-config.xml` is used much less often than in plain JSF.

In fact, once you have all the basic descriptors set up, the *only* XML you need to write as you add new functionality to a Seam application is orchestration: navigation rules or jBPM process definitions. Seam takes the view that *process flow* and *configuration data* are the only things that truly belong in XML.

In this simple example, we don't even need a navigation rule, since we decided to embed the view id in our action code.

1.2.1.7. The EJB deployment descriptor: `ejb-jar.xml`

The `ejb-jar.xml` file integrates Seam with EJB3, by attaching the `SeamInterceptor` to all session beans in the archive.

```
<ejb-jar xmlns="http://java.sun.com/xml/ns/javaee"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
        http://java.sun.com/xml/ns/javaee/ejb-jar_3_0.xsd"
    version="3.0">

    <interceptors>
        <interceptor>
            <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
        </interceptor>
    </interceptors>

    <assembly-descriptor>
        <interceptor-binding>
            <ejb-name>*</ejb-name>
            <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
        </interceptor-binding>
    </assembly-descriptor>

</ejb-jar>
```

1.2.1.8. The EJB persistence deployment descriptor: persistence.xml

The persistence.xml file tells the EJB persistence provider where to find the datasource, and contains some vendor-specific settings. In this case, enables automatic schema export at startup time.

```
<?xml version="1.0" encoding="UTF-8"?>
<persistence xmlns="http://java.sun.com/xml/ns/persistence"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/persistence
    http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd"
  version="1.0">

  <persistence-unit name="userDatabase">
    <provider>org.hibernate.ejb.HibernatePersistence</provider>
    <jta-data-source>java:/DefaultDS</jta-data-source>
    <properties>
      <property name="hibernate.hbm2ddl.auto" value="create-drop"/>
    </properties>
  </persistence-unit>

</persistence>
```

1.2.1.9. The view: register.xhtml and registered.xhtml

The view pages for a Seam application could be implemented using any technology that supports JSF. In this example we use Facelets, because we think it's better than JSP.

Example 1.7.

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
  xmlns:s="http://jboss.com/products/seam/taglib"
  xmlns:h="http://java.sun.com/jsf/html"
  xmlns:f="http://java.sun.com/jsf/core">

  <head>
    <title>Register New User</title>
  </head>
  <body>
    <f:view>
      <h:form>
        <s:validateAll>
          <h:panelGrid columns="2">
            Username: <h:inputText value="#{user.username}" required="true"/>
            Real Name: <h:inputText value="#{user.name}" required="true"/>
            Password: <h:inputSecret value="#{user.password}" required="true"/>
          </h:panelGrid>
        </s:validateAll>
        <h:messages/>
        <h:commandButton value="Register" action="#{register.register}"/>
      </h:form>
    </f:view>
  </body>
</html>
```

The only thing here that is specific to Seam is the `<s:validateAll>` tag. This JSF component tells JSF to validate all the contained input fields against the Hibernate Validator annotations specified on the entity bean.

Example 1.8.

```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:f="http://java.sun.com/jsf/core">

    <head>
        <title>Successfully Registered New User</title>
    </head>
    <body>
        <f:view>
            Welcome, #{user.name}, you are successfully registered as #{user.username}.
        </f:view>
    </body>

</html>
```

This is a boring old Facelets page using some embedded EL. There is nothing specific to Seam here.

1.2.1.10. The EAR deployment descriptor: application.xml

Finally, since our application is deployed as an EAR, we need a deployment descriptor there, too.

Example 1.9.

```
<?xml version="1.0" encoding="UTF-8"?>
<application xmlns="http://java.sun.com/xml/ns/javaee"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
        http://java.sun.com/xml/ns/javaee/application_5.xsd"
    version="5">

    <display-name>Seam Registration</display-name>

    <module>
        <web>
            <web-uri>jboss-seam-registration.war</web-uri>
            <context-root>/seam-registration</context-root>
        </web>
    </module>
    <module>
        <ejb>jboss-seam-registration.jar</ejb>
    </module>
    <module>
        <ejb>jboss-seam.jar</ejb>
    </module>
    <module>
        <java>jboss-el.jar</java>
    </module>

</application>
```

This deployment descriptor links modules in the enterprise archive and binds the web application to the context root /seam-registration.

We've now seen *all* the files in the entire application!

1.2.2. How it works

When the form is submitted, JSF asks Seam to resolve the variable named `user`. Since there is no value already bound to that name (in any Seam context), Seam instantiates the `user` component, and returns the resulting `User` entity bean instance to JSF after storing it in the Seam session context.

The form input values are now validated against the Hibernate Validator constraints specified on the `User` entity. If the constraints are violated, JSF redisplay the page. Otherwise, JSF binds the form input values to properties of the `User` entity bean.

Next, JSF asks Seam to resolve the variable named `register`. Seam finds the `RegisterAction` stateless session bean in the stateless context and returns it. JSF invokes the `register()` action listener method.

Seam intercepts the method call and injects the `User` entity from the Seam session context, before continuing the invocation.

The `register()` method checks if a user with the entered username already exists. If so, an error message is queued with the `FacesMessages` component, and a null outcome is returned, causing a page redisplay. The `FacesMessages` component interpolates the JSF expression embedded in the message string and adds a JSF `FacesMessage` to the view.

If no user with that username exists, the `"/registered.xhtml"` outcome triggers a browser redirect to the `registered.xhtml` page. When JSF comes to render the page, it asks Seam to resolve the variable named `user` and uses property values of the returned `User` entity from Seam's session scope.

1.3. Clickable lists in Seam: the messages example

Clickable lists of database search results are such an important part of any online application that Seam provides special functionality on top of JSF to make it easier to query data using EJB-QL or HQL and display it as a clickable list using a JSF `<h:dataTable>`. The messages example demonstrates this functionality.



1.3.1. Understanding the code

The message list example has one entity bean, Message, one session bean, MessageListBean and one JSP.

1.3.1.1. The entity bean: Message.java

The Message entity defines the title, text, date and time of a message, and a flag indicating whether the message has been read:

Example 1.10.

```
@Entity
@Name("message")
@Scope(EVENT)
public class Message implements Serializable
{
    private Long id;
    private String title;
    private String text;
    private boolean read;
    private Date datetime;

    @Id @GeneratedValue
    public Long getId() {
        return id;
    }
}
```

```

    }
    public void setId(Long id) {
        this.id = id;
    }

    @NotNull @Length(max=100)
    public String getTitle() {
        return title;
    }
    public void setTitle(String title) {
        this.title = title;
    }

    @NotNull @Lob
    public String getText() {
        return text;
    }
    public void setText(String text) {
        this.text = text;
    }

    @NotNull
    public boolean isRead() {
        return read;
    }
    public void setRead(boolean read) {
        this.read = read;
    }

    @NotNull
    @Basic @Temporal(TemporalType.TIMESTAMP)
    public Date getDatetime() {
        return datetime;
    }
    public void setDatetime(Date datetime) {
        this.datetime = datetime;
    }
}

```

1.3.1.2. The stateful session bean: `MessageManagerBean.java`

Just like in the previous example, we have a session bean, `MessageManagerBean`, which defines the action listener methods for the two buttons on our form. One of the buttons selects a message from the list, and displays that message. The other button deletes a message. So far, this is not so different to the previous example.

But `MessageManagerBean` is also responsible for fetching the list of messages the first time we navigate to the message list page. There are various ways the user could navigate to the page, and not all of them are preceded by a JSF action—the user might have bookmarked the page, for example. So the job of fetching the message list takes place in a Seam *factory method*, instead of in an action listener method.

We want to cache the list of messages in memory between server requests, so we will make this a stateful session bean.

Example 1.11.

```

@Stateful
@Scope(SESSION)
@Name("messageManager")
public class MessageManagerBean implements Serializable, MessageManager
{

    @DataModel

```

(1)

```

private List<Message> messageList;

@DataModelSelection                                (2)
@Out(required=false)                               (3)
private Message message;

@PersistenceContext(type=EXTENDED)                 (4)
private EntityManager em;

@Factory("messageList")                            (5)
public void findMessages()
{
    messageList = em.createQuery("from Message msg order by msg.datetime desc")
                      .getResultList();
}

public void select()                                (6)
{
    message.setRead(true);
}

public void delete()                                (7)
{
    messageList.remove(message);
    em.remove(message);
    message=null;
}

@Remove                                             (8)
public void destroy() {}
}

```

- (1) The `@DataModel` annotation exposes an attribute of type `java.util.List` to the JSF page as an instance of `javax.faces.model.DataModel`. This allows us to use the list in a JSF `<h:dataTable>` with clickable links for each row. In this case, the `DataModel` is made available in a session context variable named `messageList`.
- (2) The `@DataModelSelection` annotation tells Seam to inject the `List` element that corresponded to the clicked link.
- (3) The `@Out` annotation then exposes the selected value directly to the page. So ever time a row of the clickable list is selected, the `Message` is injected to the attribute of the stateful bean, and the subsequently *outjected* to the event context variable named `message`.
- (4) This stateful bean has an EJB3 *extended persistence context*. The messages retrieved in the query remain in the managed state as long as the bean exists, so any subsequent method calls to the stateful bean can update them without needing to make any explicit call to the `EntityManager`.
- (5) The first time we navigate to the JSP page, there will be no value in the `messageList` context variable. The `@Factory` annotation tells Seam to create an instance of `MessageManagerBean` and invoke the `findMessages()` method to initialize the value. We call `findMessages()` a *factory method* for messages.
- (6) The `select()` action listener method marks the selected `Message` as read, and updates it in the database.
- (7) The `delete()` action listener method removes the selected `Message` from the database.
- (8) All stateful session bean Seam components *must* have a method with no parameters marked `@Remove` that Seam uses to remove the stateful bean when the Seam context ends, and clean up any server-side state.

Note that this is a session-scoped Seam component. It is associated with the user login session, and all requests from a login session share the same instance of the component. (In Seam applications, we usually use session-scoped components sparingly.)

1.3.1.3. The session bean local interface: `MessageManager.java`

All session beans have a business interface, of course.

```
@Local
public interface MessageManager
{
    public void findMessages();
    public void select();
    public void delete();
    public void destroy();
}
```

From now on, we won't show local interfaces in our code examples.

Let's skip over components.xml, persistence.xml, web.xml, ejb-jar.xml, faces-config.xml and application.xml since they are much the same as the previous example, and go straight to the JSP.

1.3.1.4. The view: messages.jsp

The JSP page is a straightforward use of the JSF <h:dataTable> component. Again, nothing specific to Seam.

Example 1.12.

```
<%@ taglib uri="http://java.sun.com/jsf/html" prefix="h" %>
<%@ taglib uri="http://java.sun.com/jsf/core" prefix="f" %>
<html>
<head>
<title>Messages</title>
</head>
<body>
<f:view>
<h:form>
<h2>Message List</h2>
<h:outputText value="No messages to display"
    rendered="#{messageList.rowCount==0}"/>
<h:dataTable var="msg" value="#{messageList}"
    rendered="#{messageList.rowCount>0}">
    <h:column>
        <f:facet name="header">
            <h:outputText value="Read"/>
        </f:facet>
        <h:selectBooleanCheckbox value="#{msg.read}" disabled="true"/>
    </h:column>
    <h:column>
        <f:facet name="header">
            <h:outputText value="Title"/>
        </f:facet>
        <h:commandLink value="#{msg.title}" action="#{messageManager.select}"/>
    </h:column>
    <h:column>
        <f:facet name="header">
            <h:outputText value="Date/Time"/>
        </f:facet>
        <h:outputText value="#{msg.datetime}">
            <f:convertDateTime type="both" dateStyle="medium" timeStyle="short"/>
        </h:outputText>
    </h:column>
    <h:column>
        <h:commandButton value="Delete" action="#{messageManager.delete}"/>
    </h:column>
</h:dataTable>
<h3><h:outputText value="#{message.title}"/></h3>
<div><h:outputText value="#{message.text}"/></div>
</h:form>
</f:view>
</body>
```

```
</html>
```

1.3.2. How it works


The first time we navigate to the `messages.jsp` page, whether by a JSF postback (faces request) or a direct browser GET request (non-faces request), the page will try to resolve the `messageList` context variable. Since this context variable is not initialized, Seam will call the factory method `findMessages()`, which performs a query against the database and results in a `DataModel` being outjected. This `DataModel` provides the row data needed for rendering the `<h:dataTable>`.

When the user clicks the `<h:commandLink>`, JSF calls the `select()` action listener. Seam intercepts this call and injects the selected row data into the `message` attribute of the `messageManager` component. The action listener fires, marking the selected `Message` as read. At the end of the call, Seam outjects the selected `Message` to the context variable named `message`. Next, the EJB container commits the transaction, and the change to the `Message` is flushed to the database. Finally, the page is re-rendered, redisplaying the message list, and displaying the selected message below it.

If the user clicks the `<h:commandButton>`, JSF calls the `delete()` action listener. Seam intercepts this call and injects the selected row data into the `message` attribute of the `messageList` component. The action listener fires, removing the selected `Message` from the list, and also calling `remove()` on the `EntityManager`. At the end of the call, Seam refreshes the `messageList` context variable and clears the context variable named `message`. The EJB container commits the transaction, and deletes the `Message` from the database. Finally, the page is re-rendered, redisplaying the message list.

1.4. Seam and jBPM: the todo list example

jBPM provides sophisticated functionality for workflow and task management. To get a small taste of how jBPM integrates with Seam, we'll show you a simple "todo list" application. Since managing lists of tasks is such core functionality for jBPM, there is hardly any Java code at all in this example.



1.4.1. Understanding the code

The center of this example is the jBPM process definition. There are also two JSPs and two trivial JavaBeans (There was no reason to use session beans, since they do not access the database, or have any other transactional behavior). Let's start with the process definition:

Example 1.13.

```
<process-definition name="todo">
    <start-state name="start">                                (1)
        <transition to="todo"/>
    </start-state>

    <task-node name="todo">                                   (2)
        <task name="todo" description="#{todoList.description}"> (3)
            <assignment actor-id="#{actor.id}"/>                (4)
        </task>
        <transition to="done"/>
    </task-node>


    <end-state name="done"/>                                   (5)
</process-definition>
```

- (1) The `<start-state>` node represents the logical start of the process. When the process starts, it immediately transitions to the `todo` node.
- (2) The `<task-node>` node represents a *wait state*, where business process execution pauses, waiting for one or more tasks to be performed.
- (3) The `<task>` element defines a task to be performed by a user. Since there is only one task defined on this node, when it is complete, execution resumes, and we transition to the end state. The task gets its descrip-

tion from a Seam component named `todoList` (one of the JavaBeans).

- (4) Tasks need to be assigned to a user or group of users when they are created. In this case, the task is assigned to the current user, which we get from a built-in Seam component named `actor`. Any Seam component may be used to perform task assignment.
- (5) The `<end-state>` node defines the logical end of the business process. When execution reaches this node, the process instance is destroyed.

If we view this process definition using the process definition editor provided by JBossIDE, this is what it looks like:



This document defines our *business process* as a graph of nodes. This is the most trivial possible business process: there is one *task* to be performed, and when that task is complete, the business process ends.

The first JavaBean handles the login screen `login.jsp`. Its job is just to initialize the jBPM actor id using the `actor` component. (In a real application, it would also need to authenticate the user.)

Example 1.14.

```

@Name("login")
public class Login {

    @In
    private Actor actor;

    private String user;

    public String getUser() {
        return user;
    }

    public void setUser(String user) {
        this.user = user;
    }
}
  
```

```

    public String login()
    {
        actor.setId(user);
        return "/todo.jsp";
    }
}

```

Here we see the use of `@In` to inject the built-in Actor component.

The JSP itself is trivial:

Example 1.15.

```

<%@ taglib uri="http://java.sun.com/jsf/html" prefix="h"%>
<%@ taglib uri="http://java.sun.com/jsf/core" prefix="f"%>
<html>
<head>
<title>Login</title>
</head>
<body>
<h1>Login</h1>
<f:view>
    <h:form>
        <div>
            <h:inputText value="#{login.user}"/>
            <h:commandButton value="Login" action="#{login.login}"/>
        </div>
    </h:form>
</f:view>
</body>
</html>

```

The second JavaBean is responsible for starting business process instances, and ending tasks.

Example 1.16.

```

@Name("todoList")
public class TodoList {

    private String description;

    public String getDescription() (1)
    {
        return description;
    }

    public void setDescription(String description) {
        this.description = description;
    }

    @CreateProcess(definition="todo") (2)
    public void createTodo() {}

    @StartTask @EndTask (3)
    public void done() {}

}

```

- (1) The description property accepts user input from the JSP page, and exposes it to the process definition, allowing the task description to be set.

- (2) The Seam `@CreateProcess` annotation creates a new jBPM process instance for the named process definition.
- (3) The Seam `@StartTask` annotation starts work on a task. The `@EndTask` ends the task, and allows the business process execution to resume.

In a more realistic example, `@StartTask` and `@EndTask` would not appear on the same method, because there is usually work to be done using the application in order to complete the task.

Finally, the meat of the application is in `todo.jsp`:

Example 1.17.

```
<%@ taglib uri="http://java.sun.com/jsf/html" prefix="h" %>
<%@ taglib uri="http://java.sun.com/jsf/core" prefix="f" %>
<%@ taglib uri="http://jboss.com/products/seam/taglib" prefix="s" %>
<html>
<head>
<title>Todo List</title>
</head>
<body>
<h1>Todo List</h1>
<f:view>
  <h:form id="list">
    <div>
      <h:outputText value="There are no todo items."
        rendered="#{empty taskInstanceList}"/>
      <h:dataTable value="#{taskInstanceList}" var="task"
        rendered="#{not empty taskInstanceList}">
        <h:column>
          <f:facet name="header">
            <h:outputText value="Description"/>
          </f:facet>
          <h:inputText value="#{task.description}"/>
        </h:column>
        <h:column>
          <f:facet name="header">
            <h:outputText value="Created"/>
          </f:facet>
          <h:outputText value="#{task.taskMgmtInstance.processInstance.start}"
            <f:convertDateTime type="date"/>
          </h:outputText>
        </h:column>
        <h:column>
          <f:facet name="header">
            <h:outputText value="Priority"/>
          </f:facet>
          <h:inputText value="#{task.priority}" style="width: 30"/>
        </h:column>
        <h:column>
          <f:facet name="header">
            <h:outputText value="Due Date"/>
          </f:facet>
          <h:inputText value="#{task.dueDate}" style="width: 100">
            <f:convertDateTime type="date" dateStyle="short"/>
          </h:inputText>
        </h:column>
        <h:column>
          <s:button value="Done" action="#{todoList.done}" taskInstance="#{task}"/>
        </h:column>
      </h:dataTable>
    </div>
    <div>
      <h:messages/>
    </div>
    <div>
      <h:commandButton value="Update Items" action="update"/>
    </div>
  </h:form>
</f:view>
</body>
</html>
```

```

    </div>
  </h:form>
  <h:form id="new">
    <div>
      <h:inputText value="#{todoList.description}"/>
      <h:commandButton value="Create New Item" action="#{todoList.createTodo}"/>
    </div>
  </h:form>
</f:view>
</body>
</html>

```

Let's take this one piece at a time.

The page renders a list of tasks, which it gets from a built-in Seam component named `taskInstanceList`. The list is defined inside a JSF form.

```

<h:form id="list">
  <div>
    <h:outputText value="There are no todo items." rendered="#{empty taskInstanceList}"/>
    <h:dataTable value="#{taskInstanceList}" var="task"
      rendered="#{not empty taskInstanceList}">
      ...
    </h:dataTable>
  </div>
</h:form>

```

Each element of the list is an instance of the jBPM class `TaskInstance`. The following code simply displays the interesting properties of each task in the list. For the description, priority and due date, we use input controls, to allow the user to update these values.

```

<h:column>
  <f:facet name="header">
    <h:outputText value="Description"/>
  </f:facet>
  <h:inputText value="#{task.description}"/>
</h:column>
<h:column>
  <f:facet name="header">
    <h:outputText value="Created"/>
  </f:facet>
  <h:outputText value="#{task.taskMgmtInstance.processInstance.start}"
    <f:convertDateTime type="date"/>
  </h:outputText>
</h:column>
<h:column>
  <f:facet name="header">
    <h:outputText value="Priority"/>
  </f:facet>
  <h:inputText value="#{task.priority}" style="width: 30"/>
</h:column>
<h:column>
  <f:facet name="header">
    <h:outputText value="Due Date"/>
  </f:facet>
  <h:inputText value="#{task.dueDate}" style="width: 100">
    <f:convertDateTime type="date" dateStyle="short"/>
  </h:inputText>
</h:column>

```

This button ends the task by calling the action method annotated `@StartTask @EndTask`. It passes the task id to Seam as a request parameter:

```

<h:column>

```

```
<s:button value="Done" action="#{todoList.done}" taskInstance="#{task}"/>
</h:column>
```

(Note that this is using a Seam `<s:button>` JSF control from the `seam-ui.jar` package.)

This button is used to update the properties of the tasks. When the form is submitted, Seam and jBPM will make any changes to the tasks persistent. There is no need for any action listener method:

```
<h:commandButton value="Update Items" action="update"/>
```

A second form on the page is used to create new items, by calling the action method annotated `@CreateProcess`.

```
<h:form id="new">
  <div>
    <h:inputText value="#{todoList.description}"/>
    <h:commandButton value="Create New Item" action="#{todoList.createTodo}"/>
  </div>
</h:form>
```

There are several other files needed for the example, but they are just standard jBPM and Seam configuration and not very interesting.


1.4.2. How it works

TODO

1.5. Seam pageflow: the numberguess example

For Seam applications with relatively freeform (ad hoc) navigation, JSF/Seam navigation rules are a perfectly good way to define the page flow. For applications with a more constrained style of navigation, especially for user interfaces which are more stateful, navigation rules make it difficult to really understand the flow of the system. To understand the flow, you need to piece it together from the view pages, the actions and the navigation rules.

Seam allows you to use a jPDL process definition to define pageflow. The simple number guessing example shows how this is done.



1.5.1. Understanding the code

The example is implemented using one JavaBean, three JSP pages and a jPDL pageflow definition. Let's begin with the pageflow:

Example 1.18.

```
<pageflow-definition
  xmlns="http://jboss.com/products/seam/pageflow"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://jboss.com/products/seam/pageflow
                      http://jboss.com/products/seam/pageflow-2.0.xsd"
  name="numberGuess">

  <start-page name="displayGuess" view-id="/numberGuess.jspx">           (1)
    <redirect/>
    <transition name="guess" to="evaluateGuess">                       (2)
      <action expression="#{numberGuess.guess}"/>                     (3)
    </transition>
    <transition name="giveup" to="giveup"/>
  </start-page>

  <decision name="evaluateGuess" expression="#{numberGuess.correctGuess}"> (4)
    <transition name="true" to="win"/>
    <transition name="false" to="evaluateRemainingGuesses"/>
  </decision>

  <decision name="evaluateRemainingGuesses" expression="#{numberGuess.lastGuess}">
    <transition name="true" to="lose"/>
    <transition name="false" to="displayGuess"/>
  </decision>

  <page name="giveup" view-id="/giveup.jspx">
    <redirect/>
    <transition name="yes" to="lose"/>
    <transition name="no" to="displayGuess"/>
  </page>


  <page name="win" view-id="/win.jspx">
    <redirect/>
    <end-conversation/>
  </page>

  <page name="lose" view-id="/lose.jspx">
    <redirect/>
    <end-conversation/>
  </page>

</pageflow-definition>
```

- (1) The <page> element defines a wait state where the system displays a particular JSF view and waits for user input. The view-id is the same JSF view id used in plain JSF navigation rules. The redirect attribute tells Seam to use post-then-redirect when navigating to the page. (This results in friendly browser URLs.)
- (2) The <transition> element names a JSF outcome. The transition is triggered when a JSF action results in that outcome. Execution will then proceed to the next node of the pageflow graph, after invocation of any jBPM transition actions.
- (3) A transition <action> is just like a JSF action, except that it occurs when a jBPM transition occurs. The transition action can invoke any Seam component.
- (4) A <decision> node branches the pageflow, and determines the next node to execute by evaluating a JSF EL expression.

Here is what the pageflow looks like in the JBossIDE pageflow editor:



Now that we have seen the pageflow, it is very, very easy to understand the rest of the application!

Here is the main page of the application, numberGuess.jspx:

Example 1.19.

```

<?xml version="1.0"?>
<jsp:root xmlns:jsp="http://java.sun.com/JSP/Page"
  xmlns:h="http://java.sun.com/jsf/html"
  xmlns:f="http://java.sun.com/jsf/core"
  xmlns:s="http://jboss.com/products/seam/taglib"
  xmlns="http://www.w3.org/1999/xhtml"
  version="2.0">
  <jsp:output doctype-root-element="html"
    doctype-public="-//W3C//DTD XHTML 1.0 Transitional//EN"
    doctype-system="http://www.w3c.org/TR/xhtml1/DTD/xhtml1-transitional.dtd"/>
  <jsp:directive.page contentType="text/html"/>
  <html>
  <head>
    <title>Guess a number...</title>
    <link href="niceforms.css" rel="stylesheet" type="text/css" />
    <script language="javascript" type="text/javascript" src="niceforms.js" />
  </head>
  <body>
    <h1>Guess a number...</h1>
    <f:view>
      <h:form styleClass="niceform">

        <div>
          <h:messages globalOnly="true"/>
          <h:outputText value="Higher!"
            rendered="#{numberGuess.randomNumber gt numberGuess.currentGuess}"/>
          <h:outputText value="Lower!"
            rendered="#{numberGuess.randomNumber lt numberGuess.currentGuess}"/>
        </div>
      </h:form>
    </f:view>
  </body>
</html>

```

```

        <div>
            I'm thinking of a number between
            <h:outputText value="#{numberGuess.smallest}"/> and
            <h:outputText value="#{numberGuess.biggest}"/>. You have
            <h:outputText value="#{numberGuess.remainingGuesses}"/> guesses.
        </div>

        <div>
            Your guess:
            <h:inputText value="#{numberGuess.currentGuess}" id="inputGuess"
                required="true" size="3"
                rendered="#{(numberGuess.biggest-numberGuess.smallest) gt 20}">
                <f:validateLongRange maximum="#{numberGuess.biggest}"
                    minimum="#{numberGuess.smallest}"/>
            </h:inputText>
            <h:selectOneMenu value="#{numberGuess.currentGuess}"
                id="selectGuessMenu" required="true"
                rendered="#{(numberGuess.biggest-numberGuess.smallest) le 20 and
                    (numberGuess.biggest-numberGuess.smallest) gt 4}">
                <s:selectItems value="#{numberGuess.possibilities}" var="i" label="#{i}"/>
            </h:selectOneMenu>
            <h:selectOneRadio value="#{numberGuess.currentGuess}" id="selectGuessRadio"
                required="true"
                rendered="#{(numberGuess.biggest-numberGuess.smallest) le 4}">
                <s:selectItems value="#{numberGuess.possibilities}" var="i" label="#{i}"/>
            </h:selectOneRadio>
            <h:commandButton value="Guess" action="guess"/>
            <s:button value="Cheat" view="/confirm.jspx"/>
            <s:button value="Give up" action="giveup"/>
        </div>

        <div>
            <h:message for="inputGuess" style="color: red"/>
        </div>

    </h:form>
</f:view>
</body>
</html>
</jsp:root>

```

Notice how the command button names the guess transition instead of calling an action directly.

The win.jspx page is predictable:

Example 1.20.

```

<jsp:root xmlns:jsp="http://java.sun.com/JSP/Page"
    xmlns:h="http://java.sun.com/jsf/html"
    xmlns:f="http://java.sun.com/jsf/core"
    xmlns="http://www.w3.org/1999/xhtml"
    version="2.0">
    <jsp:output doctype-root-element="html"
        doctype-public="-//W3C//DTD XHTML 1.0 Transitional//EN"
        doctype-system="http://www.w3c.org/TR/xhtml1/DTD/xhtml1-transitional.dtd"/>
    <jsp:directive.page contentType="text/html"/>
    <html>
    <head>
        <title>You won!</title>
        <link href="niceforms.css" rel="stylesheet" type="text/css" />
    </head>
    <body>
        <h1>You won!</h1>
        <f:view>
            Yes, the answer was <h:outputText value="#{numberGuess.currentGuess}" />.

```

```

        It took you <h:outputText value="#{numberGuess.guessCount}" /> guesses.
        <h:outputText value="But you cheated, so it doesn't count!"
            rendered="#{numberGuess.cheat}"/>
        Would you like to <a href="numberGuess.seam">play again</a>?
    </f:view>
</body>
</html>
</jsp:root>

```

As is `lose.jspx` (which I can't be bothered copy/pasting). Finally, the JavaBean Seam component:

Example 1.21.

```

@Name("numberGuess")
@Scope(ScopeType.CONVERSATION)
public class NumberGuess implements Serializable {

    private int randomNumber;
    private Integer currentGuess;
    private int biggest;
    private int smallest;
    private int guessCount;
    private int maxGuesses;
    private boolean cheated;

    @Create
    public void begin()
    {
        randomNumber = new Random().nextInt(100);
        guessCount = 0;
        biggest = 100;
        smallest = 1;
    }

    public void setCurrentGuess(Integer guess)
    {
        this.currentGuess = guess;
    }

    public Integer getCurrentGuess()
    {
        return currentGuess;
    }

    public void guess()
    {
        if (currentGuess > randomNumber)
        {
            biggest = currentGuess - 1;
        }
        if (currentGuess < randomNumber)
        {
            smallest = currentGuess + 1;
        }
        guessCount++;
    }

    public boolean isCorrectGuess()
    {
        return currentGuess == randomNumber;
    }

    public int getBiggest()
    {
        return biggest;
    }
}

```

```

public int getSmallest()
{
    return smallest;
}

public int getGuessCount()
{
    return guessCount;
}

public boolean isLastGuess()
{
    return guessCount==maxGuesses;
}

public int getRemainingGuesses() {
    return maxGuesses-guessCount;
}

public void setMaxGuesses(int maxGuesses) {
    this.maxGuesses = maxGuesses;
}

public int getMaxGuesses() {
    return maxGuesses;
}

public int getRandomNumber() {
    return randomNumber;
}

public void cheated()
{
    cheated = true;
}

public boolean isCheat() {
    return cheated;
}

public List<Integer> getPossibilities()
{
    List<Integer> result = new ArrayList<Integer>();
    for(int i=smallest; i<=biggest; i++) result.add(i);
    return result;
}
}

```

- (1) The first time a JSP page asks for a numberGuess component, Seam will create a new one for it, and the @Create method will be invoked, allowing the component to initialize itself.

The pages.xml file starts a Seam *conversation* (much more about that later), and specifies the pageflow definition to use for the conversation's page flow.

Example 1.22.

```

<?xml version="1.0" encoding="UTF-8"?>
<pages xmlns="http://jboss.com/products/seam/pages"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://jboss.com/products/seam/pages http://jboss.com/products/seam/pages-2.0.xsd">

    <page view-id="/numberGuess.jspx">
        <begin-conversation join="true" pageflow="numberGuess"/>
    </page>

```

```
<page view-id="/confirm.jspx">
  <begin-conversation nested="true" pageflow="cheat"/>
</page>

</pages>
```

As you can see, this Seam component is pure business logic! It doesn't need to know anything at all about the user interaction flow. This makes the component potentially more reusable.

1.5.2. How it works

TODO


1.6. A complete Seam application: the Hotel Booking example

1.6.1. Introduction

The booking application is a complete hotel room reservation system incorporating the following features:

- User registration
- Login
- Logout
- Set password
- Hotel search
- Hotel selection
- Room reservation
- Reservation confirmation
- Existing reservation list

jboss suites
seam framework demo
Welcome Gavin King | Search | Settings | Logout



State management in Seam

State in Seam is *contextual*. When you click "Find Hotels", the application retrieves a list of hotels from the database and caches it in the session context. When you navigate to one of the hotel records by clicking the "View Hotel" link, a *conversation* begins. The conversation is attached to a particular tab, in a particular browser window. You can navigate to multiple hotels using "open in new tab" or "open in new window" in your web browser. Each window will execute in the context of a different conversation. The application keeps state associated with your hotel booking in the conversation context, which ensures that the concurrent conversations do not interfere with each other.

[How does the search page work?](#)

Thank you, Gavin King, your confirmation number for Doubletree is 1

Search Hotels

Maximum results:

Name	Address	City, State	Zip	Action
Marriott Courtyard	Tower Place, Buckhead	Atlanta, GA, USA	30305	View Hotel
Doubletree	Tower Place, Buckhead	Atlanta, GA, USA	30305	View Hotel
Ritz Carlton	Peachtree Rd, Buckhead	Atlanta, GA, USA	30326	View Hotel

Current Hotel Bookings

Name	Address	City, State	Check in date	Check out date	Confirmation number	Action
Doubletree	Tower Place, Buckhead	Atlanta, GA	Apr 16, 2006	Apr 17, 2006	1	Cancel

Created with JBoss EJB 3.0, Seam, MyFaces, and Facelets

The booking application uses JSF, EJB 3.0 and Seam, together with Facelets for the view. There is also a port of this application to JSF, Facelets, Seam, JavaBeans and Hibernate3.

One of the things you'll notice if you play with this application for long enough is that it is extremely *robust*. You can play with back buttons and browser refresh and opening multiple windows and entering nonsensical data as much as you like and you will find it very difficult to make the application crash. You might think that we spent weeks testing and fixing bugs to achieve this. Actually, this is not the case. Seam was designed to make it very straightforward to build robust web applications and a lot of robustness that you are probably used to having to code yourself comes naturally and automatically with Seam.

As you browse the sourcecode of the example application, and learn how the application works, observe how the declarative state management and integrated validation has been used to achieve this robustness.

1.6.2. Overview of the booking example

The project structure is identical to the previous one, to install and deploy this application, please refer to Section 1.1, “Try the examples”. Once you’ve successfully started the application, you can access it by pointing your browser to `http://localhost:8080/seam-booking/` [`http://localhost:8080/seam-booking/`]

Just nine classes (plus six session beans local interfaces) were used to implement this application. Six session bean action listeners contain all the business logic for the listed features.

- `BookingListAction` retrieves existing bookings for the currently logged in user.
- `ChangePasswordAction` updates the password of the currently logged in user.
- `HotelBookingAction` implements the core functionality of the application: hotel room searching, selection, booking and booking confirmation. This functionality is implemented as a *conversation*, so this is the most interesting class in the application.
- `RegisterAction` registers a new system user.

Three entity beans implement the application's persistent domain model.

- `Hotel` is an entity bean that represent a hotel
- `Booking` is an entity bean that represents an existing booking
- `User` is an entity bean to represents a user who can make hotel bookings

1.6.3. Understanding Seam conversations

We encourage you browse the sourcecode at your pleasure. In this tutorial we'll concentrate upon one particular piece of functionality: hotel search, selection, booking and confirmation. From the point of view of the user, everything from selecting a hotel to confirming a booking is one continuous unit of work, a *conversation*. Searching, however, is *not* part of the conversation. The user can select multiple hotels from the same search results page, in different browser tabs.

Most web application architectures have no first class construct to represent a conversation. This causes enormous problems managing state associated with the conversation. Usually, Java web applications use a combination of two techniques: first, some state is thrown into the `HttpSession`; second, persistable state is flushed to the database after every request, and reconstructed from the database at the beginning of each new request.

Since the database is the least scalable tier, this often results in an utterly unacceptable lack of scalability. Added latency is also a problem, due to the extra traffic to and from the database on every request. To reduce this redundant traffic, Java applications often introduce a data (second-level) cache that keeps commonly accessed data between requests. This cache is necessarily inefficient, because invalidation is based upon an LRU policy instead of being based upon when the user has finished working with the data. Furthermore, because the cache is shared between many concurrent transactions, we've introduced a whole raft of problem's associated with keeping the cached state consistent with the database.

Now consider the state held in the `HttpSession`. By very careful programming, we might be able to control the size of the session data. This is a lot more difficult than it sounds, since web browsers permit ad hoc non-linear navigation. But suppose we suddenly discover a system requirement that says that a user is allowed to have *multiple concurrent conversations*, halfway through the development of the system (this has happened to me). Developing mechanisms to isolate session state associated with different concurrent conversations, and incorporat-

ing failsafes to ensure that conversation state is destroyed when the user aborts one of the conversations by closing a browser window or tab is not for the faint hearted (I've implemented this stuff twice so far, once for a client application, once for Seam, but I'm famously psychotic).

Now there is a better way.

Seam introduces the *conversation context* as a first class construct. You can safely keep conversational state in this context, and be assured that it will have a well-defined lifecycle. Even better, you won't need to be continually pushing data back and forth between the application server and the database, since the conversation context is a natural cache of data that the user is currently working with.

Usually, the components we keep in the conversation context are stateful session beans. (We can also keep entity beans and JavaBeans in the conversation context.) There is an ancient canard in the Java community that stateful session beans are a scalability killer. This may have been true in 1998 when WebFoobar 1.0 was released. It is no longer true today. Application servers like JBoss AS have extremely sophisticated mechanisms for stateful session bean state replication. (For example, the JBoss EJB3 container performs fine-grained replication, replicating only those bean attribute values which actually changed.) Note that all the traditional technical arguments for why stateful beans are inefficient apply equally to the `HttpSession`, so the practice of shifting state from business tier stateful session bean components to the web session to try and improve performance is unbelievably misguided. It is certainly possible to write unscalable applications using stateful session beans, by using stateful beans incorrectly, or by using them for the wrong thing. But that doesn't mean you should *never* use them. Anyway, Seam guides you toward a safe usage model. Welcome to 2005.

OK, I'll stop ranting now, and get back to the tutorial.

The booking example application shows how stateful components with different scopes can collaborate together to achieve complex behaviors. The main page of the booking application allows the user to search for hotels. The search results are kept in the Seam session scope. When the user navigates to one of these hotels, a conversation begins, and a conversation scoped component calls back to the session scoped component to retrieve the selected hotel.

The booking example also demonstrates the use of RichFaces Ajax to implement rich client behavior without the use of handwritten JavaScript.

The search functionality is implemented using a session-scope stateful session bean, similar to the one we saw in the message list example above.

Example 1.23.

```
@Stateful (1)
@Name("hotelSearch")
@Scope(ScopeType.SESSION)
@Restrict("#{identity.loggedIn}") (2)
public class HotelSearchingAction implements HotelSearching
{
    @PersistenceContext
    private EntityManager em;

    private String searchString;
    private int pageSize = 10;
    private int page;

    @DataModel (3)
    private List<Hotel> hotels;

    public void find()
    {
        page = 0;
    }
}
```

```

        queryHotels();
    }
    public void nextPage()
    {
        page++;
        queryHotels();
    }

    private void queryHotels()
    {
        hotels =
            em.createQuery("select h from Hotel h where lower(h.name) like #{pattern} " +
                           "or lower(h.city) like #{pattern} " +
                           "or lower(h.zip) like #{pattern} " +
                           "or lower(h.address) like #{pattern}")
                .setMaxResults(pageSize)
                .setFirstResult( page * pageSize )
                .getResultList();
    }

    public boolean isNextPageAvailable()
    {
        return hotels!=null && hotels.size()==pageSize;
    }

    public int getPageSize() {
        return pageSize;
    }

    public void setPageSize(int pageSize) {
        this.pageSize = pageSize;
    }

    @Factory(value="pattern", scope=ScopeType.EVENT)
    public String getSearchPattern()
    {
        return searchString==null ?
            "%" : '%' + searchString.toLowerCase().replace('*', '%') + '%';
    }

    public String getSearchString()
    {
        return searchString;
    }

    public void setSearchString(String searchString)
    {
        this.searchString = searchString;
    }

    @Remove
    public void destroy() {}
}

```

(4)

- (1) The EJB standard `@Stateful` annotation identifies this class as a stateful session bean. Stateful session beans are scoped to the conversation context by default.
- (2) The `@Restrict` annotation applies a security restriction to the component. It restricts access to the component allowing only logged-in users. The security chapter explains more about security in Seam.
- (3) The `@DataModel` annotation exposes a `List` as a JSF `ListDataModel`. This makes it easy to implement clickable lists for search screens. In this case, the list of hotels is exposed to the page as a `ListDataModel` in the conversation variable named `hotels`.
- (4) The EJB standard `@Remove` annotation specifies that a stateful session bean should be removed and its state destroyed after invocation of the annotated method. In Seam, all stateful session beans must define a method with no parameters marked `@Remove`. This method will be called when Seam destroys the session context.

The main page of the application is a Facelets page. Let's look at the fragment which relates to searching for hotels:

Example 1.24.

```
<div class="section">

    <span class="errors">
        <h:messages globalOnly="true"/>
    </span>

    <h1>Search Hotels</h1>

    <h:form id="searchCriteria">
        <fieldset>
            <h:inputText id="searchString" value="#{hotelSearch.searchString}"
                style="width: 165px;">
                <a:support event="onkeyup" actionListener="#{hotelSearch.find}"
                    reRender="searchResults" /> (1)
            </h:inputText>
            &#160;
            <a:commandButton id="findHotels" value="Find Hotels" action="#{hotelSearch.find}"
                reRender="searchResults"/>
            &#160;
            <a:status> (2)
                <f:facet name="start">
                    <h:graphicImage value="/img/spinner.gif"/>
                </f:facet>
            </a:status>
            <br/>
            <h:outputLabel for="pageSize">Maximum results:</h:outputLabel>&#160;
            <h:selectOneMenu value="#{hotelSearch.pageSize}" id="pageSize">
                <f:selectItem itemLabel="5" itemValue="5"/>
                <f:selectItem itemLabel="10" itemValue="10"/>
                <f:selectItem itemLabel="20" itemValue="20"/>
            </h:selectOneMenu>
        </fieldset>
    </h:form>

</div>

<a:outputPanel id="searchResults"> (3)
    <div class="section">
        <h:outputText value="No Hotels Found"
            rendered="#{hotels != null and hotels.rowCount==0}"/>
        <h:dataTable id="hotels" value="#{hotels}" var="hot"
            rendered="#{hotels.rowCount>0}">
            <h:column>
                <f:facet name="header">Name</f:facet>
                #{hot.name}
            </h:column>
            <h:column>
                <f:facet name="header">Address</f:facet>
                #{hot.address}
            </h:column>
            <h:column>
                <f:facet name="header">City, State</f:facet>
                #{hot.city}, #{hot.state}, #{hot.country}
            </h:column>
            <h:column>
                <f:facet name="header">Zip</f:facet>
                #{hot.zip}
            </h:column>
            <h:column>
                <f:facet name="header">Action</f:facet>
                <s:link id="viewHotel" value="View Hotel" (4)
                    action="#{hotelBooking.selectHotel(hot)}/>
            </h:column>
        </h:dataTable>
    </div>
</a:outputPanel>
</div>
```

```

    </h:dataTable>
    <s:link value="More results" action="#{hotelSearch.nextPage}"
          rendered="#{hotelSearch.nextPageAvailable}"/>
  </div>
</a:outputPanel>

```

- (1) The RichFaces Ajax `<a:support>` tag allows a JSF action event listener to be called by asynchronous XMLHttpRequest when a JavaScript event like `onkeyup` occurs. Even better, the `reRender` attribute lets us render a fragment of the JSF page and perform a partial page update when the asynchronous response is received.
- (2) The RichFaces Ajax `<a:status>` tag lets us display a cheesy animated image while we wait for asynchronous requests to return.
- (3) The RichFaces Ajax `<a:outputPanel>` tag defines a region of the page which can be re-rendered by an asynchronous request.
- (4) The Seam `<s:link>` tag lets us attach a JSF action listener to an ordinary (non-JavaScript) HTML link. The advantage of this over the standard JSF `<h:commandLink>` is that it preserves the operation of "open in new window" and "open in new tab". Also notice that we use a method binding with a parameter: `#{hotelBooking.selectHotel(hot)}`. This is not possible in the standard Unified EL, but Seam provides an extension to the EL that lets you use parameters on any method binding expression.

If you're wondering how navigation occurs, you can find all the rules in `WEB-INF/pages.xml`; this is discussed in Section 5.1.1.2, "Navigation".

This page displays the search results dynamically as we type, and lets us choose a hotel and pass it to the `selectHotel()` method of the `HotelBookingAction`, which is where the *really* interesting stuff is going to happen.

Now let's see how the booking example application uses a conversation-scoped stateful session bean to achieve a natural cache of persistent data related to the conversation. The following code example is pretty long. But if you think of it as a list of scripted actions that implement the various steps of the conversation, it's understandable. Read the class from top to bottom, as if it were a story.

Example 1.25.

```

@Stateful
@Name("hotelBooking")
@Restrict("#{identity.loggedIn}")
public class HotelBookingAction implements HotelBooking
{
    @PersistenceContext(type=EXTENDED)                                (1)
    private EntityManager em;

    @In
    private User user;

    @In(required=false) @Out
    private Hotel hotel;

    @In(required=false)
    @Out(required=false)                                             (2)
    private Booking booking;

    @In
    private FacesMessages facesMessages;

    @In
    private Events events;

```

```

@Logger
private Log log;

private boolean bookingValid;

@Begin (3)
public void selectHotel(Hotel selectedHotel)
{
    hotel = em.merge(selectedHotel);
}

public void bookHotel()
{
    booking = new Booking(hotel, user);
    Calendar calendar = Calendar.getInstance();
    booking.setCheckinDate( calendar.getTime() );
    calendar.add(Calendar.DAY_OF_MONTH, 1);
    booking.setCheckoutDate( calendar.getTime() );
}

public void setBookingDetails()
{
    Calendar calendar = Calendar.getInstance();
    calendar.add(Calendar.DAY_OF_MONTH, -1);
    if ( booking.getCheckinDate().before( calendar.getTime() ) )
    {
        facesMessages.addToControl("checkinDate", "Check in date must be a future date");
        bookingValid=false;
    }
    else if ( !booking.getCheckinDate().before( booking.getCheckoutDate() ) )
    {
        facesMessages.addToControl("checkoutDate",
                                   "Check out date must be later than check in date");
        bookingValid=false;
    }
    else
    {
        bookingValid=true;
    }
}

public boolean isBookingValid()
{
    return bookingValid;
}

@end (4)
public void confirm()
{
    em.persist(booking);
    facesMessages.add("Thank you, #{user.name}, your confirmation number " +
                      "for #{hotel.name} is #{booking.id}");
    log.info("New booking: #{booking.id} for #{user.username}");
    events.raiseTransactionSuccessEvent("bookingConfirmed");
}

@end
public void cancel() {}

@Remove (5)
public void destroy() {}

```

- (1) This bean uses an EJB3 *extended persistence context*, so that any entity instances remain managed for the whole lifecycle of the stateful session bean.
- (2) The @Out annotation declares that an attribute value is *outjected* to a context variable after method invocations. In this case, the context variable named hotel will be set to the value of the hotel instance variable after every action listener invocation completes.
- (3) The @Begin annotation specifies that the annotated method begins a *long-running conversation*, so the

current conversation context will not be destroyed at the end of the request. Instead, it will be reassociated with every request from the current window, and destroyed either by timeout due to conversation inactivity or invocation of a matching `@End` method.

- (4) The `@End` annotation specifies that the annotated method ends the current long-running conversation, so the current conversation context will be destroyed at the end of the request.
- (5) This EJB remove method will be called when Seam destroys the conversation context. Don't forget to define this method!

`HotelBookingAction` contains all the action listener methods that implement selection, booking and booking confirmation, and holds state related to this work in its instance variables. We think you'll agree that this code is much cleaner and simpler than getting and setting `HttpSession` attributes.

Even better, a user can have multiple isolated conversations per login session. Try it! Log in, run a search, and navigate to different hotel pages in multiple browser tabs. You'll be able to work on creating two different hotel reservations at the same time. If you leave any one conversation inactive for long enough, Seam will eventually time out that conversation and destroy its state. If, after ending a conversation, you backbutton to a page of that conversation and try to perform an action, Seam will detect that the conversation was already ended, and redirect you to the search page.

1.6.4. The Seam UI control library

If you check inside the WAR file for the booking application, you'll find `seam-ui.jar` in the `WEB-INF/lib` directory. This package contains a number of JSF custom controls that integrate with Seam. The booking application uses the `<s:link>` control for navigation from the search screen to the hotel page:

```
<s:link value="View Hotel" action="#{hotelBooking.selectHotel(hot)}"/>
```

The use of `<s:link>` here allows us to attach an action listener to a HTML link without breaking the browser's "open in new window" feature. The standard JSF `<h:commandLink>` does not work with "open in new window". We'll see later that `<s:link>` also offers a number of other useful features, including conversation propagation rules.

The booking application uses some other Seam and RichFaces Ajax controls, especially on the `/book.xhtml` page. We won't get into the details of those controls here, but if you want to understand this code, please refer to the chapter covering Seam's functionality for JSF form validation.

1.6.5. The Seam Debug Page

The WAR also includes `seam-debug.jar`. The Seam debug page will be available if this jar is deployed in `WEB-INF/lib`, along with the Facelets, and if you set the debug property of the `init` component:

```
<core:init jndi-pattern="@jndiPattern@" debug="true"/>
```

This page lets you browse and inspect the Seam components in any of the Seam contexts associated with your current login session. Just point your browser at `http://localhost:8080/seam-booking/debug.seam` [`http://localhost:8080/seam-booking/debug.seam`].

JBoss Seam Debug Page

This page allows you to view and inspect any component in any Seam context associated with the current session.

Conversations

conversation id	activity	description	view id	
4	1:51:34 AM - 1:51:34 AM	Search hotels: M	/main.xhtml	Select conversation context
6	1:51:40 AM - 1:52:23 AM	Book hotel: Marriott Courtyard	/book.xhtml	Select conversation context

- Component (booking)

checkinDate	Fri Jan 20 20:52:20 EST 2006
checkoutDate	Sat Jan 21 20:52:20 EST 2006
class	class org.jboss.seam.example.booking.Booking
creditCard	
description	Marriott Courtyard, Jan 20, 2006 to Jan 21, 2006
hotel	Hotel(Tower Place, Buckhead,Atlanta,30305)
id	
user	User(gavin)

- Conversation Context (6)

booking
conversation
hotel
hotelBooking
hotels

- Business Process Context

Empty business process context

+ Session Context

+ Application Context

1.7. A complete application featuring Seam and jBPM: the DVD Store example

The DVD Store demo application shows the practical usage of jBPM for both task management and pageflow.

The user screens take advantage of a jPDL pageflow to implement searching and shopping cart functionality.

JBoss Seam DVD Store Demo

[Search for Movies](#)
[My Orders](#)

Search Results

Add to cart	Title	Actor	Price
☐	Life is Beautiful	Roberto Benini	\$12.00
☐	Finding Nemo	Albert Brooks	\$22.49
☐	March of the Penguins	Morgan Freeman	\$16.98
☐	Indiana Jones and the Temple of Doom	Harrison Ford	\$19.99
☐	Clear and Present Danger	Harrison Ford	\$19.99
☐	Roman Holiday	Audrey Hepburn	\$12.99
☐	Breakfast at Tiffany's	Audrey Hepburn	\$12.99
☐	Sabrina	Audrey Hepburn	\$12.99
☐	Sabrina	Harrison Ford	\$19.99
☐	Kill Bill Vol. 1	Uma Thurman	\$19.99
☐	Kill Bill Vol. 2	Uma Thurman	\$19.99
☐	Lost in Translation	Bill Murray	\$19.99
☐	Broken Flowers	Bill Murray	\$19.99
☐	Better Off Dead	John Cusak	\$8.99
☐	Grosse Pointe Blank	John Cusak	\$11.99
☐	High Fidelity	John Cusak	\$14.99
☐	Somewhere in Time	Christopher Reeve	\$11.24
☐	Superman - The Movie	Christopher Reeve	\$14.99
☐	Superman II	Christopher Reeve	\$14.99
☐	Superman III	Christopher Reeve	\$14.99

Update Shopping Cart

Welcome, Harry

Thank you for choosing the DVD Store

Logout

Search for DVDs:

Title:

Actor:

Category:

Any

Results Per Page:

20

Search

Shopping Cart

1 Napoleon Dynamite

Total:\$14.06

Checkout

Done

The administration screens take use jBPM to manage the approval and shipping cycle for orders. The business process may even be changed dynamically, by selecting a different process definition!

JBoss Seam DVD Store Demo

Manage Orders

Order Management

Pending orders are shown here on the order management screen for the store manager to process. Rather than being data-driven, order management is process-driven. A JBoss jBPM process assigns fulfillment tasks to the manager based on the version of the process loaded. The manager can change the version of the process at any time using the admin options box to the right.

- Order process 1 sends orders immediately to shipping, where the manager should ship the order and record the tracking number for the user to see.
- Order process 2 adds an approval step where the manager is first given the chance to approve the order before sending it to shipping. In each case, the status of the order is shown in the customer's order list.
- Order process 3 introduces a decision node. Only orders over \$100.00 need to be accepted. Smaller orders are automatically approved for shipping.

Task Assignment

Order Id	Order Amount	Customer	Task	
5	\$12.99	user1	ship	Assign
7	\$77.70	user2	ship	Assign

Order Acceptance

There are no orders to be accepted.

Shipping

Order Id	Order Amount	Customer	
6	\$94.95	user1	Ship

Done

Welcome, Albus

Thank you for choosing the DVD Store

Statistics

Inventory
28 sold, 2473 in stock

Sales
\$437.63 from 7 orders

Admin Options

Process Management

ordermanagement3

TODO

Look in the dvdstore directory.

1.8. An example of Seam with Hibernate: the Hibernate Booking example

The Hibernate Booking demo is a straight port of the Booking demo to an alternative architecture that uses Hibernate for persistence and JavaBeans instead of session beans.

TODO

Look in the hibernate directory.

1.9. A RESTful Seam application: the Blog example

Seam makes it very easy to implement applications which keep state on the server-side. However, server-side state is not always appropriate, especially in for functionality that serves up *content*. For this kind of problem

we often need to let the user bookmark pages and have a relatively stateless server, so that any page can be accessed at any time, via the bookmark. The Blog example shows how to implement a RESTful application using Seam. Every page of the application can be bookmarked, including the search results page.



The Blog example demonstrates the use of "pull"-style MVC, where instead of using action listener methods to retrieve data and prepare the data for the view, the view pulls data from components as it is being rendered.

1.9.1. Using "pull"-style MVC

This snippet from the `index.xhtml` facelets page displays a list of recent blog entries:

Example 1.26.

```
<h:dataTable value="#{blog.recentBlogEntries}" var="blogEntry" rows="3">
  <h:column>
    <div class="blogEntry">
      <h3>#{blogEntry.title}</h3>
      <div>
        <h:outputText escape="false"
          value="#{blogEntry.excerpt==null ? blogEntry.body : blogEntry.excerpt}"/>
      </div>
      <p>
        <h:outputLink value="entry.seam" rendered="#{blogEntry.excerpt!=null}">
          <f:param name="blogEntryId" value="#{blogEntry.id}"/>
          Read more...
        </h:outputLink>
      </p>
    </h:column>
  </h:dataTable>
```

```

        <p>
            [Posted on
            <h:outputText value="#{blogEntry.date}">
                <f:convertDateTime timeZone="#{blog.timeZone}"
                                   locale="#{blog.locale}" type="both"/>
            </h:outputText>]
            &#160;
            <h:outputLink value="entry.seam">[Link]
                <f:param name="blogEntryId" value="#{blogEntry.id}"/>
            </h:outputLink>
        </p>
    </div>
</h:column>
</h:dataTable>

```

If we navigate to this page from a bookmark, how does the data used by the `<h:dataTable>` actually get initialized? Well, what happens is that the `Blog` is retrieved lazily—"pulled"—when needed, by a Seam component named `blog`. This is the opposite flow of control to what is usual in traditional web action-based frameworks like Struts.

Example 1.27.

```

@Name("blog")
@Scope(ScopeType.STATELESS)
@AutoCreate
public class BlogService
{
    @In EntityManager entityManager;                                (1)

    @Unwrap                                                         (2)
    public Blog getBlog()
    {
        return (Blog) entityManager.createQuery("select distinct b from Blog b left join fetch b.blogEnt
            .setHint("org.hibernate.cacheable", true)
            .getSingleResult();
    }
}

```

- (1) This component uses a *seam-managed persistence context*. Unlike the other examples we've seen, this persistence context is managed by Seam, instead of by the EJB3 container. The persistence context spans the entire web request, allowing us to avoid any exceptions that occur when accessing unfetched associations in the view.
- (2) The `@Unwrap` annotation tells Seam to provide the return value of the method—the `Blog`—instead of the actual `BlogService` component to clients. This is the *Seam manager component pattern*.

This is good so far, but what about bookmarking the result of form submissions, such as a search results page?

1.9.2. Bookmarkable search results page

The blog example has a tiny form in the top right of each page that allows the user to search for blog entries. This is defined in a file, `menu.xhtml`, included by the facelets template, `template.xhtml`:

Example 1.28.

```
<div id="search">
  <h:form>
    <h:inputText value="#{searchAction.searchPattern}"/>
    <h:commandButton value="Search" action="/search.xhtml"/>
  </h:form>
</div>
```

To implement a bookmarkable search results page, we need to perform a browser redirect after processing the search form submission. Because we used the JSF view id as the action outcome, Seam automatically redirects to the view id when the form is submitted. Alternatively, we could have defined a navigation rule like this:

Example 1.29.

```
<navigation-rule>
  <navigation-case>
    <from-outcome>searchResults</from-outcome>
    <to-view-id>/search.xhtml</to-view-id>
    <redirect/>
  </navigation-case>
</navigation-rule>
```

Then the form would have looked like this:

Example 1.30.

```
<div id="search">
  <h:form>
    <h:inputText value="#{searchAction.searchPattern}"/>
    <h:commandButton value="Search" action="searchResults"/>
  </h:form>
</div>
```

But when we redirect, we need to include the values submitted with the form as request parameters, to get a bookmarkable URL like `http://localhost:8080/seam-blog/search.seam?searchPattern=seam`. JSF does not provide an easy way to do this, but Seam does. We use a *Seam page parameter*, defined in `WEB-INF/pages.xml`:

Example 1.31.

```
<pages>
  <page view-id="/search.xhtml">
    <param name="searchPattern" value="#{searchService.searchPattern}"/>
  </page>
  ...
</pages>
```

This tells Seam to include the value of `#{searchService.searchPattern}` as a request parameter named `searchPattern` when redirecting to the page, and then re-apply the value of that parameter to the model before rendering the page.

The redirect takes us to the `search.xhtml` page:

Example 1.32.

```

<h:dataTable value="#{searchResults}" var="blogEntry">
  <h:column>
    <div>
      <h:outputLink value="entry.seam">
        <f:param name="blogEntryId" value="#{blogEntry.id}"/>
        #{blogEntry.title}
      </h:outputLink>
      posted on
      <h:outputText value="#{blogEntry.date}">
        <f:convertDateTime timeZone="#{blog.timeZone}" locale="#{blog.locale}" type="both"/>
      </h:outputText>
    </div>
  </h:column>
</h:dataTable>

```

Which again uses "pull"-style MVC to retrieve the actual search results:

Example 1.33.

```

@Name("searchService")
public class SearchService
{
    @In
    private EntityManager entityManager;

    private String searchPattern;

    @Factory("searchResults")
    public List<BlogEntry> getSearchResults()
    {
        if (searchPattern==null)
        {
            return null;
        }
        else
        {
            return entityManager.createQuery("select be from BlogEntry be " +
                "where lower(be.title) like :searchPattern " +
                "lower(be.body) like :searchPattern order by be.date desc")
                .setParameter( "searchPattern", getSqlSearchPattern() )
                .setMaxResults(100)
                .getResultList();
        }
    }

    private String getSqlSearchPattern()
    {
        return searchPattern==null ? "" :
            '%' + searchPattern.toLowerCase().replace('*', '%').replace('?', '_') + '%';
    }

    public String getSearchPattern()
    {
        return searchPattern;
    }

    public void setSearchPattern(String searchPattern)
    {
        this.searchPattern = searchPattern;
    }
}

```

1.9.3. Using "push"-style MVC in a RESTful application

Very occasionally, it makes more sense to use push-style MVC for processing RESTful pages, and so Seam provides the notion of a *page action*. The Blog example uses a page action for the blog entry page, `entry.xhtml`. Note that this is a little bit contrived, it would have been easier to use pull-style MVC here as well.

The `entryAction` component works much like an action class in a traditional push-MVC action-oriented framework like Struts:

Example 1.34.

```
@Name("entryAction")
@Scope(STATELESS)
public class EntryAction
{
    @In(create=true)
    private Blog blog;

    @Out
    private BlogEntry blogEntry;

    public void loadBlogEntry(String id) throws EntryNotFoundException
    {
        blogEntry = blog.getBlogEntry(id);
        if (blogEntry==null) throw new EntryNotFoundException(id);
    }
}
```

Page actions are also declared in `pages.xml`:

Example 1.35.

```
<pages>
...

<page view-id="/entry.xhtml" action="#{entryAction.loadBlogEntry(blogEntry.id)}">
    <param name="blogEntryId" value="#{blogEntry.id}"/>
</page>

<page view-id="/post.xhtml" action="#{loginAction.challenge}"/>

<page view-id="*" action="#{blog.hitCount.hit}"/>
</pages>
```

Notice that the example is using page actions for some other functionality—the login challenge, and the pageview counter. Also notice the use of a parameter in the page action method binding. This is not a standard feature of JSF EL, but Seam lets you use it, not just for page actions, but also in JSF method bindings.

When the `entry.xhtml` page is requested, Seam first binds the page parameter `blogEntryId` to the model, then runs the page action, which retrieves the needed data—the `blogEntry`—and places it in the Seam event context. Finally, the following is rendered:

Example 1.36.

```

<div class="blogEntry">
  <h3>#{blogEntry.title}</h3>
  <div>
    <h:outputText escape="false" value="#{blogEntry.body}"/>
  </div>
  <p>
    [Posted on#{160};
    <h:outputText value="#{blogEntry.date}">
      <f:convertDateTime timezone="#{blog.timeZone}"
        locale="#{blog.locale}" type="both"/>
    </h:outputText>]
  </p>
</div>

```

If the blog entry is not found in the database, the `EntryNotFoundException` exception is thrown. We want this exception to result in a 404 error, not a 505, so we annotate the exception class:

Example 1.37.

```

@ApplicationException(rollback=true)
@HttpError(errorCode=HttpServletResponse.SC_NOT_FOUND)
public class EntryNotFoundException extends Exception
{
    EntryNotFoundException(String id)
    {
        super("entry not found: " + id);
    }
}

```

An alternative implementation of the example does not use the parameter in the method binding:

Example 1.38.

```

@Name("entryAction")
@Scope(STATELESS)
public class EntryAction
{
    @In(create=true)
    private Blog blog;

    @In @Out
    private BlogEntry blogEntry;

    public void loadBlogEntry() throws EntryNotFoundException
    {
        blogEntry = blog.getBlogEntry( blogEntry.getId() );
        if (blogEntry==null) throw new EntryNotFoundException(id);
    }
}

```

```

<pages>
  ...

  <page view-id="/entry.xhtml" action="#{entryAction.loadBlogEntry}">
    <param name="blogEntryId" value="#{blogEntry.id}"/>
  </page>

  ...
</pages>

```

It is a matter of taste which implementation you prefer.

Chapter 2. Getting started with Seam, using seam-gen

The Seam distribution includes a command line utility that makes it really easy to set up an Eclipse project, generate some simple Seam skeleton code, and reverse engineer an application from a preexisting database.

This is the easy way to get your feet wet with Seam, and gives you some ammunition for next time you find yourself trapped in an elevator with one of those tedious Ruby guys ranting about how great and wonderful his new toy is for building totally trivial applications that put things in databases.

In this release, seam-gen works best for people with JBoss AS. You can use the generated project with other J2EE or Java EE 5 application servers by making a few manual changes to the project configuration.

You *can* use seam-gen without Eclipse, but in this tutorial, we want to show you how to use it in conjunction with Eclipse for debugging and integration testing. If you don't want to install Eclipse, you can still follow along with this tutorial—all steps can be performed from the command line.

Seam-gen is basically just a big ugly Ant script wrapped around Hibernate Tools, together with some templates. That makes it easy to customize if you need to.

2.1. Before you start

Make sure you have JDK 5 or JDK 6, JBoss AS 4.2 and Ant 1.6, along with recent versions of Eclipse, the JBoss IDE plugin for Eclipse and the TestNG plugin for Eclipse correctly installed before starting. Add your JBoss installation to the JBoss Server View in Eclipse. Start JBoss in debug mode. Finally, start a command prompt in the directory where you unzipped the Seam distribution.

JBoss has sophisticated support for hot re-deployment of WARs and EARs. Unfortunately, due to bugs in the JVM, repeated redeployment of an EAR—which is common during development—eventually causes the JVM to run out of perm gen space. For this reason, we recommend running JBoss in a JVM with a large perm gen space at development time. If you're running JBoss from JBoss IDE, you can configure this in the server launch configuration, under "VM arguments". We suggest the following values:

```
-Xms512m -Xmx1024m -XX:PermSize=256m -XX:MaxPermSize=512
```

If you don't have so much memory available, the following is our minimum recommendation:

```
-Xms256m -Xmx512m -XX:PermSize=128m -XX:MaxPermSize=256
```

If you're running JBoss from the command line, you can configure the JVM options in `bin/run.conf`.

If you don't want to bother with this stuff now, you don't have to—come back to it later, when you get your first `OutOfMemoryException`.

2.2. Setting up a new Eclipse project

The first thing we need to do is configure seam-gen for your environment: JBoss AS installation directory, Eclipse workspace, and database connection. It's easy, just type:

```
cd jboss-seam-2.0.x
seam setup
```

And you will be prompted for the needed information:

```
~/workspace/jboss-seam$ ./seam setup
Buildfile: build.xml

init:

setup:
  [echo] Welcome to seam-gen :-)
  [input] Enter your Java project workspace (the directory that contains your Seam projects) [C:/Program Files/Java/j2se6.0_03]
/Users/pmuir/workspace
  [input] Enter your JBoss home directory [C:/Program Files/jboss-4.2.2.GA] [C:/Program Files/jboss-4.2.2.GA]
/Applications/jboss-4.2.2.GA
  [input] Enter the project name [myproject] [myproject]
helloworld
  [echo] Accepted project name as: helloworld
  [input] Select a RichFaces skin (not applicable if using ICEfaces) [blueSky] ([blueSky], classic,
  [input] Is this project deployed as an EAR (with EJB components) or a WAR (with no EJB support) [ear]
  [input] Enter the Java package name for your session beans [com.mydomain.helloworld] [com.mydomain.helloworld]
org.jboss.helloworld
  [input] Enter the Java package name for your entity beans [org.jboss.helloworld] [org.jboss.helloworld]
  [input] Enter the Java package name for your test cases [org.jboss.helloworld.test] [org.jboss.helloworld.test]
  [input] What kind of database are you using? [hsqldb] ([hsqldb], mysql, oracle, postgres, mssql, db2, h2, hsql)
mysql
  [input] Enter the Hibernate dialect for your database [org.hibernate.dialect.MySQLDialect] [org.hibernate.dialect.MySQLDialect]
  [input] Enter the filesystem path to the JDBC driver jar [lib/hsqldb.jar] [lib/hsqldb.jar]
/Users/pmuir/java/mysql.jar
  [input] Enter JDBC driver class for your database [com.mysql.jdbc.Driver] [com.mysql.jdbc.Driver]
  [input] Enter the JDBC URL for your database [jdbc:mysql:///test] [jdbc:mysql:///test]
jdbc:mysql:///helloworld
  [input] Enter database username [sa] [sa]
pmuir
  [input] Enter database password [] []
  [input] skipping input as property hibernate.default_schema.new has already been set.
  [input] Enter the database catalog name (it is OK to leave this blank) [] []
  [input] Are you working with tables that already exist in the database? [n] (y, [n], )
y
  [input] Do you want to drop and recreate the database tables and data in import.sql each time you
n
  [input] Enter your ICEfaces home directory (leave blank to omit ICEfaces) [] []

[propertyfile] Creating new property file: /Users/pmuir/workspace/jboss-seam/seam-gen/build.properties
  [echo] Installing JDBC driver jar to JBoss server
  [echo] Type 'seam create-project' to create the new project

BUILD SUCCESSFUL
Total time: 1 minute 32 seconds
~/workspace/jboss-seam $
```

The tool provides sensible defaults, which you can accept by just pressing enter at the prompt.

The most important choice you need to make is between EAR deployment and WAR deployment of your project. EAR projects support EJB 3.0 and require Java EE 5. WAR projects do not support EJB 3.0, but may be deployed to a J2EE environment. The packaging of a WAR is also simpler to understand. If you installed an EJB3-ready application server like JBoss, choose ear. Otherwise, choose war. We'll assume that you've chosen an EAR deployment for the rest of the tutorial, but you can follow exactly the same steps for a WAR deployment.

If you are working with an existing data model, make sure you tell seam-gen that the tables already exist in the

database.

The settings are stored in `seam-gen/build.properties`, but you can also modify them simply by running `seam` setup a second time.

Now we can create a new project in our Eclipse workspace directory, by typing:

```
seam new-project
```

```
C:\Projects\jboss-seam>seam new-project
Buildfile: build.xml

...

new-project:
  [echo] A new Seam project named 'helloworld' was created in the C:\Projects directory
  [echo] Type 'seam explode' and go to http://localhost:8080/helloworld
  [echo] Eclipse Users: Add the project into Eclipse using File > New > Project and select General
  [echo] NetBeans Users: Open the project in NetBeans

BUILD SUCCESSFUL
Total time: 7 seconds
C:\Projects\jboss-seam>
```

This copies the Seam jars, dependent jars and the JDBC driver jar to a new Eclipse project, and generates all needed resources and configuration files, a facelets template file and stylesheet, along with Eclipse metadata and an Ant build script. The Eclipse project will be automatically deployed to an exploded directory structure in JBoss AS as soon as you add the project using `New -> Project... -> General -> Project -> Next`, typing the Project name (helloworld in this case), and then clicking `Finish`. Do not select `Java Project` from the New Project wizard.

If your default JDK in Eclipse is not a Java SE 5 or Java SE 6 JDK, you will need to select a Java SE 5 compliant JDK using `Project -> Properties -> Java Compiler`.

Alternatively, you can deploy the project from outside Eclipse by typing `seam explode`.

Go to `http://localhost:8080/helloworld` to see a welcome page. This is a facelets page, `view/home.xhtml`, using the template `view/layout/template.xhtml`. You can edit this page, or the template, in eclipse, and see the results *immediately*, by clicking `refresh` in your browser.

Don't get scared by the XML configuration documents that were generated into the project directory. They are mostly standard Java EE stuff, the stuff you need to create once and then never look at again, and they are 90% the same between all Seam projects. (They are so easy to write that even seam-gen can do it.)

The generated project includes three database and persistence configurations. The `jboss-beans.xml`, `persistence-test.xml` and `import-test.sql` files are used when running the TestNG unit tests against HSQLDB. The database schema and the test data in `import-test.sql` is always exported to the database before running tests. The `myproject-dev-ds.xml`, `persistence-dev.xml` and `import-dev.sql` files are for use when deploying the application to your development database. The schema might be exported automatically at deployment, depending upon whether you told seam-gen that you are working with an existing database. The `myproject-prod-ds.xml`, `persistence-prod.xml` and `import-prod.sql` files are for use when deploying the application to your production database. The schema is not exported automatically at deployment.

2.3. Creating a new action

If you're used to traditional action-style web frameworks, you're probably wondering how you can create a

simple web page with a stateless action method in Java. If you type:

```
seam new-action
```

Seam will prompt for some information, and generate a new facelets page and Seam component for your project.

```
C:\Projects\jboss-seam>seam new-action
Buildfile: build.xml

validate-workspace:

validate-project:

action-input:
  [input] Enter the Seam component name
ping
  [input] Enter the local interface name [Ping]

  [input] Enter the bean class name [PingBean]

  [input] Enter the action method name [ping]

  [input] Enter the page name [ping]

setup-filters:

new-action:
  [echo] Creating a new stateless session bean component with an action method
  [copy] Copying 1 file to C:\Projects\helloworld\src\action\org\jboss\helloworld
  [copy] Copying 1 file to C:\Projects\helloworld\src\action\org\jboss\helloworld
  [copy] Copying 1 file to C:\Projects\helloworld\src\action\org\jboss\helloworld\test
  [copy] Copying 1 file to C:\Projects\helloworld\src\action\org\jboss\helloworld\test
  [copy] Copying 1 file to C:\Projects\helloworld\view
  [echo] Type 'seam restart' and go to http://localhost:8080/helloworld/ping.seam

BUILD SUCCESSFUL
Total time: 13 seconds
C:\Projects\jboss-seam>
```

Because we've added a new Seam component, we need to restart the exploded directory deployment. You can do this by typing `seam restart`, or by running the `restart` target in the generated project `build.xml` file from inside Eclipse. Another way to force a restart is to edit the file `resources/META-INF/application.xml` in Eclipse. *Note that you do not need to restart JBoss each time you change the application.*

Now go to `http://localhost:8080/helloworld/ping.seam` and click the button. You can see the code behind this action by looking in the project `src` directory. Put a breakpoint in the `ping()` method, and click the button again.

Finally, locate the `PingTest.xml` file in the test package and run the integration tests using the TestNG plugin for Eclipse. Alternatively, run the tests using `seam test` or the `test` target of the generated build.

2.4. Creating a form with an action

The next step is to create a form. Type:

```
seam new-form
```

```
C:\Projects\jboss-seam>seam new-form
Buildfile: C:\Projects\jboss-seam\seam-gen\build.xml
```

```

validate-workspace:

validate-project:

action-input:
    [input] Enter the Seam component name
hello
    [input] Enter the local interface name [Hello]

    [input] Enter the bean class name [HelloBean]

    [input] Enter the action method name [hello]

    [input] Enter the page name [hello]

setup-filters:

new-form:
    [echo] Creating a new stateful session bean component with an action method
    [copy] Copying 1 file to C:\Projects\hello\src\com\hello
    [copy] Copying 1 file to C:\Projects\hello\src\com\hello
    [copy] Copying 1 file to C:\Projects\hello\src\com\hello\test
    [copy] Copying 1 file to C:\Projects\hello\view
    [copy] Copying 1 file to C:\Projects\hello\src\com\hello\test
    [echo] Type 'seam restart' and go to http://localhost:8080/hello/hello.seam

BUILD SUCCESSFUL
Total time: 5 seconds
C:\Projects\jboss-seam>

```

Restart the application again, and go to <http://localhost:8080/helloworld/hello.seam>. Then take a look at the generated code. Run the test. Try adding some new fields to the form and Seam component (remember to restart the deployment each time you change the Java code).

2.5. Generating an application from an existing database

Manually create some tables in your database. (If you need to switch to a different database, just run seam setup again.) Now type:

```
seam generate-entities
```

Restart the deployment, and go to <http://localhost:8080/helloworld>. You can browse the database, edit existing objects, and create new objects. If you look at the generated code, you'll probably be amazed how simple it is! Seam was designed so that data access code is easy to write by hand, even for people who don't want to cheat by using seam-gen.

2.6. Generating an application from existing JPA/EJB3 entities

Place your existing, valid entity classes inside the `src/model`. Now type

```
seam generate-ui
```

Restart the deployment, and go to <http://localhost:8080/helloworld>.

2.7. Deploying the application as an EAR

Finally, we want to be able to deploy the application using standard Java EE 5 packaging. First, we need to remove the exploded directory by running `seam unexplode`. To deploy the EAR, we can type `seam deploy` at the command prompt, or run the `deploy` target of the generated project build script. You can undeploy using `seam undeploy` or the `undeploy` target.

By default, the application will be deployed with the *dev profile*. The EAR will include the `persistence-dev.xml` and `import-dev.sql` files, and the `myproject-dev-ds.xml` file will be deployed. You can change the profile, and use the *prod profile*, by typing

```
seam -Dprofile=prod deploy
```

You can even define new deployment profiles for your application. Just add appropriately named files to your project—for example, `persistence-staging.xml`, `import-staging.sql` and `myproject-staging-ds.xml`—and select the name of the profile using `-Dprofile=staging`.

2.8. Seam and incremental hot deployment

When you deploy your Seam application as an exploded directory, you'll get some support for incremental hot deployment at development time. You need to enable debug mode in both Seam and Facelets, by adding this line to `components.xml`:

```
<core:init debug="true"/>
```

Now, the following files may be redeployed without requiring a full restart of the web application:

- any facelets page
- any `pages.xml` file

But if we want to change any Java code, we still need to do a full restart of the application. (In JBoss this may be accomplished by touching the top level deployment descriptor: `application.xml` for an EAR deployment, or `web.xml` for a WAR deployment.)

But if you really want a fast edit/compile/test cycle, Seam supports incremental redeployment of JavaBean components. To make use of this functionality, you must deploy the JavaBean components into the `WEB-INF/dev` directory, so that they will be loaded by a special Seam classloader, instead of by the WAR or EAR classloader.

You need to be aware of the following limitations:

- the components must be JavaBean components, they cannot be EJB3 beans (we are working on fixing this limitation)
- entities can never be hot-deployed
- components deployed via `components.xml` may not be hot-deployed
- the hot-deployable components will not be visible to any classes deployed outside of `WEB-INF/dev`
- Seam debug mode must be enabled

If you create a WAR project using seam-gen, incremental hot deployment is available out of the box for classes

in the `src/action` source directory. However, seam-gen does not support incremental hot deployment for EAR projects.

2.9. Using Seam with JBoss 4.0

Seam 2.0 was developed for JavaServer Faces 1.2. When using JBoss AS, we recommend using JBoss 4.2, which bundles the JSF 1.2 reference implementation. However, it is still possible to use Seam 2.0 on the JBoss 4.0 platform. There are two basic steps required to do this: install an EJB3-enabled version of JBoss 4.0 and replace MyFaces with the JSF 1.2 reference implementation. Once you complete these steps, Seam 2.0 applications can be deployed to JBoss 4.0.

2.9.1. Install JBoss 4.0

JBoss 4.0 does not ship a default configuration compatible with Seam. To run Seam, you must install JBoss 4.0.5 using the JEMS 1.2 installer with the `ejb3` profile selected. Seam will not run with an installation that doesn't include EJB3 support. The JEMS installer can be downloaded from <http://labs.jboss.com/jemsinstaller/downloads>.

2.9.2. Install the JSF 1.2 RI

The web configuration for JBoss 4.0 can be found in the `server/default/deploy/jbossweb-tomcat55.sar`. You'll need to delete `myfaces-api.jar` and `myfaces-impl.jar` from the `jsf-libs` directory. Then, you'll need to copy `jsf-api.jar`, `jsf-impl.jar`, `el-api.jar`, and `el-ri.jar` to that directory. The JSF JARs can be found in the `Seam lib` directory. The `el` JARs can be obtained from the Seam 1.2 release.

You'll also need to edit the `conf/web.xml`, replacing `myfaces-impl.jar` with `jsf-impl.jar`.

Chapter 3. The contextual component model

The two core concepts in Seam are the notion of a *context* and the notion of a *component*. Components are stateful objects, usually EJBs, and an instance of a component is associated with a context, and given a name in that context. *Bijection* provides a mechanism for aliasing internal component names (instance variables) to contextual names, allowing component trees to be dynamically assembled, and reassembled by Seam.

Let's start by describing the contexts built in to Seam.

3.1. Seam contexts

Seam contexts are created and destroyed by the framework. The application does not control context demarcation via explicit Java API calls. Context are usually implicit. In some cases, however, contexts are demarcated via annotations.

The basic Seam contexts are:

- Stateless context
- Event (or request) context
- Page context
- Conversation context
- Session context
- Business process context
- Application context

You will recognize some of these contexts from servlet and related specifications. However, two of them might be new to you: *conversation context*, and *business process context*. One reason state management in web applications is so fragile and error-prone is that the three built-in contexts (request, session and application) are not especially meaningful from the point of view of the business logic. A user login session, for example, is a fairly arbitrary construct in terms of the actual application work flow. Therefore, most Seam components are scoped to the conversation or business process contexts, since they are the contexts which are most meaningful in terms of the application.

Let's look at each context in turn.

3.1.1. Stateless context

Components which are truly stateless (stateless session beans, primarily) always live in the stateless context (this is really a non-context). Stateless components are not very interesting, and are arguably not very object-oriented. Nevertheless, they are important and often useful.

3.1.2. Event context

The event context is the "narrowest" stateful context, and is a generalization of the notion of the web request context to cover other kinds of events. Nevertheless, the event context associated with the lifecycle of a JSF re-

quest is the most important example of an event context, and the one you will work with most often. Components associated with the event context are destroyed at the end of the request, but their state is available and well-defined for at least the lifecycle of the request.

When you invoke a Seam component via RMI, or Seam Remoting, the event context is created and destroyed just for the invocation.

3.1.3. Page context

The page context allows you to associate state with a particular instance of a rendered page. You can initialize state in your event listener, or while actually rendering the page, and then have access to it from any event that originates from that page. This is especially useful for functionality like clickable lists, where the list is backed by changing data on the server side. The state is actually serialized to the client, so this construct is extremely robust with respect to multi-window operation and the back button.

3.1.4. Conversation context

The conversation context is a truly central concept in Seam. A *conversation* is a unit of work from the point of view of the user. It might span several interactions with the user, several requests, and several database transactions. But to the user, a conversation solves a single problem. For example, "book hotel", "approve contract", "create order" are all conversations. You might like to think of a conversation implementing a single "use case" or "user story", but the relationship is not necessarily quite exact.

A conversation holds state associated with "what the user is doing now, in this window". A single user may have multiple conversations in progress at any point in time, usually in multiple windows. The conversation context allows us to ensure that state from the different conversations does not collide and cause bugs.

It might take you some time to get used to thinking of applications in terms of conversations. But once you get used to it, we think you'll love the notion, and never be able to not think in terms of conversations again!

Some conversations last for just a single request. Conversations that span multiple requests must be demarcated using annotations provided by Seam.

Some conversations are also *tasks*. A task is a conversation that is significant in terms of a long-running business process, and has the potential to trigger a business process state transition when it is successfully completed. Seam provides a special set of annotations for task demarcation.

Conversations may be *nested*, with one conversation taking place "inside" a wider conversation. This is an advanced feature.

Usually, conversation state is actually held by Seam in the servlet session between requests. Seam implements configurable *conversation timeout*, automatically destroying inactive conversations, and thus ensuring that the state held by a single user login session does not grow without bound if the user abandons conversations.

Seam serializes processing of concurrent requests that take place in the same long-running conversation context, in the same process.

Alternatively, Seam may be configured to keep conversational state in the client browser.

3.1.5. Session context

A session context holds state associated with the user login session. While there are some cases where it is use-

ful to share state between several conversations, we generally frown on the use of session context for holding components other than global information about the logged in user.

In a JSR-168 portal environment, the session context represents the portlet session.

3.1.6. Business process context

The business process context holds state associated with the long running business process. This state is managed and made persistent by the BPM engine (JBoss jBPM). The business process spans multiple interactions with multiple users, so this state is shared between multiple users, but in a well-defined manner. The current task determines the current business process instance, and the lifecycle of the business process is defined externally using a *process definition language*, so there are no special annotations for business process demarcation.

3.1.7. Application context

The application context is the familiar servlet context from the servlet spec. Application context is mainly useful for holding static information such as configuration data, reference data or metamodels. For example, Seam stores its own configuration and metamodel in the application context.

3.1.8. Context variables

A context defines a namespace, a set of *context variables*. These work much the same as session or request attributes in the servlet spec. You may bind any value you like to a context variable, but usually we bind Seam component instances to context variables.

So, within a context, a component instance is identified by the context variable name (this is usually, but not always, the same as the component name). You may programatically access a named component instance in a particular scope via the Contexts class, which provides access to several thread-bound instances of the Context interface:

```
User user = (User) Contexts.getSessionContext().get("user");
```

You may also set or change the value associated with a name:

```
Contexts.getSessionContext().set("user", user);
```

Usually, however, we obtain components from a context via injection, and put component instances into a context via outjection.

3.1.9. Context search priority

Sometimes, as above, component instances are obtained from a particular known scope. Other times, all stateful scopes are searched, in *priority order*. The order is as follows:

- Event context
- Page context
- Conversation context

- Session context
- Business process context
- Application context

You can perform a priority search by calling `Contexts.lookupInStatefulContexts()`. Whenever you access a component by name from a JSF page, a priority search occurs.

3.1.10. Concurrency model

Neither the servlet nor EJB specifications define any facilities for managing concurrent requests originating from the same client. The servlet container simply lets all threads run concurrently and leaves enforcing thread-safeness to application code. The EJB container allows stateless components to be accessed concurrently, and throws an exception if multiple threads access a stateful session bean.

This behavior might have been okay in old-style web applications which were based around fine-grained, synchronous requests. But for modern applications which make heavy use of many fine-grained, asynchronous (AJAX) requests, concurrency is a fact of life, and must be supported by the programming model. Seam weaves a concurrency management layer into its context model.

The Seam session and application contexts are multithreaded. Seam will allow concurrent requests in a context to be processed concurrently. The event and page contexts are by nature single threaded. The business process context is strictly speaking multi-threaded, but in practice concurrency is sufficiently rare that this fact may be disregarded most of the time. Finally, Seam enforces a *single thread per conversation per process* model for the conversation context by serializing concurrent requests in the same long-running conversation context.

Since the session context is multithreaded, and often contains volatile state, session scope components are always protected by Seam from concurrent access. Seam serializes requests to session scope session beans and JavaBeans by default (and detects and breaks any deadlocks that occur). This is not the default behaviour for application scoped components however, since application scoped components do not usually hold volatile state and because synchronization at the global level is *extremely* expensive. However, you can force a serialized threading model on any session bean or JavaBean component by adding the `@Synchronized` annotation.

This concurrency model means that AJAX clients can safely use volatile session and conversational state, without the need for any special work on the part of the developer.

3.2. Seam components

Seam components are POJOs (Plain Old Java Objects). In particular, they are JavaBeans or EJB 3.0 enterprise beans. While Seam does not require that components be EJBs and can even be used without an EJB 3.0 compliant container, Seam was designed with EJB 3.0 in mind and includes deep integration with EJB 3.0. Seam supports the following *component types*.

- EJB 3.0 stateless session beans
- EJB 3.0 stateful session beans
- EJB 3.0 entity beans
- JavaBeans

- EJB 3.0 message-driven beans

3.2.1. Stateless session beans

Stateless session bean components are not able to hold state across multiple invocations. Therefore, they usually work by operating upon the state of other components in the various Seam contexts. They may be used as JSF action listeners, but cannot provide properties to JSF components for display.

Stateless session beans always live in the stateless context.

Stateless session beans can be accessed concurrently as a new instance is used for each request. Assigning the instance to the request is the responsibility of the EJB3 container (normally instances will be allocated from a reusable pool meaning that you may find any instance variables contain data from previous uses of the bean).

Stateless session beans are the least interesting kind of Seam component.

Seam stateless session bean components may be instantiated using `Component.getInstance()` or `@In(create=true)`. They should not be directly instantiated via JNDI lookup or the new operator.

3.2.2. Stateful session beans

Stateful session bean components are able to hold state not only across multiple invocations of the bean, but also across multiple requests. Application state that does not belong in the database should usually be held by stateful session beans. This is a major difference between Seam and many other web application frameworks. Instead of sticking information about the current conversation directly in the `HttpSession`, you should keep it in instance variables of a stateful session bean that is bound to the conversation context. This allows Seam to manage the lifecycle of this state for you, and ensure that there are no collisions between state relating to different concurrent conversations.

Stateful session beans are often used as JSF action listener, and as backing beans that provide properties to JSF components for display or form submission.

By default, stateful session beans are bound to the conversation context. They may never be bound to the page or stateless contexts.

Concurrent requests to session-scoped stateful session beans are always serialized by Seam.

Seam stateful session bean components may be instantiated using `Component.getInstance()` or `@In(create=true)`. They should not be directly instantiated via JNDI lookup or the new operator.

3.2.3. Entity beans

Entity beans may be bound to a context variable and function as a seam component. Because entities have a persistent identity in addition to their contextual identity, entity instances are usually bound explicitly in Java code, rather than being instantiated implicitly by Seam.

Entity bean components do not support bijection or context demarcation. Nor does invocation of an entity bean trigger validation.

Entity beans are not usually used as JSF action listeners, but do often function as backing beans that provide properties to JSF components for display or form submission. In particular, it is common to use an entity as a backing bean, together with a stateless session bean action listener to implement create/update/delete type func-

tionality.

By default, entity beans are bound to the conversation context. They may never be bound to the stateless context.

Note that it in a clustered environment is somewhat less efficient to bind an entity bean directly to a conversation or session scoped Seam context variable than it would be to hold a reference to the entity bean in a stateful session bean. For this reason, not all Seam applications define entity beans to be Seam components.

Seam entity bean components may be instantiated using `Component.getInstance()`, `@In(create=true)` or directly using the `new` operator.

3.2.4. JavaBeans

JavaBeans may be used just like a stateless or stateful session bean. However, they do not provide the functionality of a session bean (declarative transaction demarcation, declarative security, efficient clustered state replication, EJB 3.0 persistence, timeout methods, etc).

In a later chapter, we show you how to use Seam and Hibernate without an EJB container. In this use case, components are JavaBeans instead of session beans. Note, however, that in many application servers it is somewhat less efficient to cluster conversation or session scoped Seam JavaBean components than it is to cluster stateful session bean components.

By default, JavaBeans are bound to the event context.

Concurrent requests to session-scoped JavaBeans are always serialized by Seam.

Seam JavaBean components may be instantiated using `Component.getInstance()` or `@In(create=true)`. They should not be directly instantiated using the `new` operator.

3.2.5. Message-driven beans

Message-driven beans may function as a seam component. However, message-driven beans are called quite differently to other Seam components - instead of invoking them via the context variable, they listen for messages sent to a JMS queue or topic.

Message-driven beans may not be bound to a Seam context. Nor do they have access to the session or conversation state of their "caller". However, they do support bijection and some other Seam functionality.

Message-driven beans are never instantiated by the application. They are instantiated by the EJB container when a message is received.

3.2.6. Interception

In order to perform its magic (bijection, context demarcation, validation, etc), Seam must intercept component invocations. For JavaBeans, Seam is in full control of instantiation of the component, and no special configuration is needed. For entity beans, interception is not required since bijection and context demarcation are not defined. For session beans, we must register an EJB interceptor for the session bean component. We could use an annotation, as follows:

```
@Stateless
@Interceptors(SeamInterceptor.class)
public class LoginAction implements Login {
    ...
}
```

```
}
```

But a much better way is to define the interceptor in `ejb-jar.xml`.

```
<interceptors>
  <interceptor>
    <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
  </interceptor>
</interceptors>

<assembly-descriptor>
  <interceptor-binding>
    <ejb-name>*</ejb-name>
    <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
  </interceptor-binding>
</assembly-descriptor>
```

3.2.7. Component names

All seam components need a name. We can assign a name to a component using the `@Name` annotation:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    ...
}
```

This name is the *seam component name* and is not related to any other name defined by the EJB specification. However, seam component names work just like JSF managed bean names and you can think of the two concepts as identical.

`@Name` is not the only way to define a component name, but we always need to specify the name *somewhere*. If we don't, then none of the other Seam annotations will function.

Just like in JSF, a seam component instance is usually bound to a context variable with the same name as the component name. So, for example, we would access the `LoginAction` using `Contexts.getStatelessContext().get("loginAction")`. In particular, whenever Seam itself instantiates a component, it binds the new instance to a variable with the component name. However, again like JSF, it is possible for the application to bind a component to some other context variable by programmatic API call. This is only useful if a particular component serves more than one role in the system. For example, the currently logged in user might be bound to the `currentUser` session context variable, while a user that is the subject of some administration functionality might be bound to the `user conversation` context variable.

For very large applications, and for built-in seam components, qualified names are often used.

```
@Name("com.jboss.myapp.loginAction")
@Stateless
public class LoginAction implements Login {
    ...
}
```

We may use the qualified component name both in Java code and in JSF's expression language:

```
<h:commandButton type="submit" value="Login"
  action="#{com.jboss.myapp.loginAction.login}"/>
```

Since this is noisy, Seam also provides a means of aliasing a qualified name to a simple name. Add a line like this to the `components.xml` file:

```
<factory name="loginAction" scope="STATELESS" value="#{com.jboss.myapp.loginAction}"/>
```

All of the built-in Seam components have qualified names, but most of them are aliased to a simple name by the `components.xml` file included in the Seam jar.

3.2.8. Defining the component scope

We can override the default scope (context) of a component using the `@Scope` annotation. This lets us define what context a component instance is bound to, when it is instantiated by Seam.

```
@Name("user")
@Entity
@Scope(SESSION)
public class User {
    ...
}
```

`org.jboss.seam.ScopeType` defines an enumeration of possible scopes.

3.2.9. Components with multiple roles

Some Seam component classes can fulfill more than one role in the system. For example, we often have a `User` class which is usually used as a session-scoped component representing the current user but is used in user administration screens as a conversation-scoped component. The `@Role` annotation lets us define an additional named role for a component, with a different scope—it lets us bind the same component class to different context variables. (Any Seam component *instance* may be bound to multiple context variables, but this lets us do it at the class level, and take advantage of auto-instantiation.)

```
@Name("user")
@Entity
@Scope(CONVERSATION)
@Role(name="currentUser", scope=SESSION)
public class User {
    ...
}
```

The `@Roles` annotation lets us specify as many additional roles as we like.

```
@Name("user")
@Entity
@Scope(CONVERSATION)
@Roles({@Role(name="currentUser", scope=SESSION),
        @Role(name="tempUser", scope=EVENT)})
public class User {
    ...
}
```

3.2.10. Built-in components

Like many good frameworks, Seam eats its own dogfood and is implemented mostly as a set of built-in Seam interceptors (see later) and Seam components. This makes it easy for applications to interact with built-in components at runtime or even customize the basic functionality of Seam by replacing the built-in components with custom implementations. The built-in components are defined in the Seam namespace `org.jboss.seam.core` and the Java package of the same name.

The built-in components may be injected, just like any Seam components, but they also provide convenient

static instance() methods:

```
FacesMessages.instance().add("Welcome back, #{user.name}!");
```

3.3. Bijection

Dependency injection or *inversion of control* is by now a familiar concept to most Java developers. Dependency injection allows a component to obtain a reference to another component by having the container "inject" the other component to a setter method or instance variable. In all dependency injection implementations that we have seen, injection occurs when the component is constructed, and the reference does not subsequently change for the lifetime of the component instance. For stateless components, this is reasonable. From the point of view of a client, all instances of a particular stateless component are interchangeable. On the other hand, Seam emphasizes the use of stateful components. So traditional dependency injection is no longer a very useful construct. Seam introduces the notion of *bijection* as a generalization of injection. In contrast to injection, bijection is:

- *contextual* - bijection is used to assemble stateful components from various different contexts (a component from a "wider" context may even have a reference to a component from a "narrower" context)
- *bidirectional* - values are injected from context variables into attributes of the component being invoked, and also *outjected* from the component attributes back out to the context, allowing the component being invoked to manipulate the values of contextual variables simply by setting its own instance variables
- *dynamic* - since the value of contextual variables changes over time, and since Seam components are stateful, bijection takes place every time a component is invoked

In essence, bijection lets you alias a context variable to a component instance variable, by specifying that the value of the instance variable is injected, outjected, or both. Of course, we use annotations to enable bijection.

The `@In` annotation specifies that a value should be injected, either into an instance variable:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    @In User user;
    ...
}
```

or into a setter method:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    User user;

    @In
    public void setUser(User user) {
        this.user=user;
    }

    ...
}
```

By default, Seam will do a priority search of all contexts, using the name of the property or instance variable that is being injected. You may wish to specify the context variable name explicitly, using, for example, `@In("currentUser")`.

If you want Seam to create an instance of the component when there is no existing component instance bound to the named context variable, you should specify `@In(create=true)`. If the value is optional (it can be null), specify `@In(required=false)`.

For some components, it can be repetitive to have to specify `@In(create=true)` everywhere they are used. In such cases, you can annotate the component `@AutoCreate`, and then it will always be created, whenever needed, even without the explicit use of `create=true`.

You can even inject the value of an expression:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    @In("#{user.username}") String username;
    ...
}
```

Injected values are disinjected (i.e, set to null) immediately after method completion and outjection.

(There is much more information about component lifecycle and injection in the next chapter.)

The `@Out` annotation specifies that an attribute should be outjected, either from an instance variable:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    @Out User user;
    ...
}
```

or from a getter method:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    User user;

    @Out
    public User getUser() {
        return user;
    }

    ...
}
```

An attribute may be both injected and outjected:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    @In @Out User user;
    ...
}
```

or:

```
@Name("loginAction")
@Stateless
public class LoginAction implements Login {
    User user;

    @In
    public void setUser(User user) {
```

```
        this.user=user;
    }

    @Out
    public User getUser() {
        return user;
    }

    ...
}
```

3.4. Lifecycle methods

Session bean and entity bean Seam components support all the usual EJB 3.0 lifecycle callback (`@PostConstruct`, `@PreDestroy`, etc). But Seam also supports the use of any of these callbacks with JavaBean components. However, since these annotations are not available in a J2EE environment, Seam defines two additional component lifecycle callbacks, equivalent to `@PostConstruct` and `@PreDestroy`.

The `@Create` method is called after Seam instantiates a component. Components may define only one `@Create` method.

The `@Destroy` method is called when the context that the Seam component is bound to ends. Components may define only one `@Destroy` method.

In addition, stateful session bean components *must* define a method with no parameters annotated `@Remove`. This method is called by Seam when the context ends.

Finally, a related annotation is the `@Startup` annotation, which may be applied to any application or session scoped component. The `@Startup` annotation tells Seam to instantiate the component immediately, when the context begins, instead of waiting until it is first referenced by a client. It is possible to control the order of instantiation of startup components by specifying `@Startup(depends={ . . . })`.

3.5. Conditional installation

The `@Install` annotation lets you control conditional installation of components that are required in some deployment scenarios and not in others. This is useful if:

- You want to mock out some infrastructural component in tests.
- You want change the implementation of a component in certain deployment scenarios.
- You want to install some components only if their dependencies are available (useful for framework authors).

`@Install` works by letting you specify *precedence* and *dependencies*.

The precedence of a component is a number that Seam uses to decide which component to install when there are multiple classes with the same component name in the classpath. Seam will choose the component with the higher precedence. There are some predefined precedence values (in ascending order):

1. `BUILT_IN` — the lowest precedence components are the components built in to Seam.
2. `FRAMEWORK` — components defined by third-party frameworks may override built-in components, but are

overridden by application components.

3. APPLICATION — the default precedence. This is appropriate for most application components.
4. DEPLOYMENT — for application components which are deployment-specific.
5. MOCK — for mock objects used in testing.

Suppose we have a component named `messageSender` that talks to a JMS queue.

```
@Name("messageSender")
public class MessageSender {
    public void sendMessage() {
        //do something with JMS
    }
}
```

In our unit tests, we don't have a JMS queue available, so we would like to stub out this method. We'll create a *mock* component that exists in the classpath when unit tests are running, but is never deployed with the application:

```
@Name("messageSender")
@Install(precedence=MOCK)
public class MockMessageSender extends MessageSender {
    public void sendMessage() {
        //do nothing!
    }
}
```

The precedence helps Seam decide which version to use when it finds both components in the classpath.

This is nice if we are able to control exactly which classes are in the classpath. But if I'm writing a reusable framework with many dependencies, I don't want to have to break that framework across many jars. I want to be able to decide which components to install depending upon what other components are installed, and upon what classes are available in the classpath. The `@Install` annotation also controls this functionality. Seam uses this mechanism internally to enable conditional installation of many of the built-in components. However, you probably won't need to use it in your application.

3.6. Logging

Who is not totally fed up with seeing noisy code like this?

```
private static final Log log = LogFactory.getLog(CreateOrderAction.class);

public Order createOrder(User user, Product product, int quantity) {
    if ( log.isDebugEnabled() ) {
        log.debug("Creating new order for user: " + user.username() +
            " product: " + product.name()
            + " quantity: " + quantity);
    }
    return new Order(user, product, quantity);
}
```

It is difficult to imagine how the code for a simple log message could possibly be more verbose. There is more lines of code tied up in logging than in the actual business logic! I remain totally astonished that the Java community has not come up with anything better in 10 years.

Seam provides a logging API that simplifies this code significantly:

```
@Logger private Log log;

public Order createOrder(User user, Product product, int quantity) {
    log.debug("Creating new order for user: #0 product: #1 quantity: #2", user.username(), product.name(), quantity);
    return new Order(user, product, quantity);
}
```

It doesn't matter if you declare the `log` variable static or not—it will work either way, except for entity bean components which require the `log` variable to be static.

Note that we don't need the noisy `if ( log.isDebugEnabled() )` guard, since string concatenation happens *inside* the `debug()` method. Note also that we don't usually need to specify the log category explicitly, since Seam knows what component it is injecting the `Log` into.

If `User` and `Product` are Seam components available in the current contexts, it gets even better:

```
@Logger private Log log;

public Order createOrder(User user, Product product, int quantity) {
    log.debug("Creating new order for user: #{user.username} product: #{product.name} quantity: #0", user.username(), product.name(), quantity);
    return new Order(user, product, quantity);
}
```

Seam logging automatically chooses whether to send output to log4j or JDK logging. If log4j is in the classpath, Seam will use it. If it is not, Seam will use JDK logging.

3.7. The `Mutable` interface and `@ReadOnly`

Many application servers feature an amazingly broken implementation of `HttpSession` clustering, where changes to the state of mutable objects bound to the session are only replicated when the application calls `setAttribute()` explicitly. This is a source of bugs that can not effectively be tested for at development time, since they will only manifest when failover occurs. Furthermore, the actual replication message contains the entire serialized object graph bound to the session attribute, which is inefficient.

Of course, EJB stateful session beans must perform automatic dirty checking and replication of mutable state and a sophisticated EJB container can introduce optimizations such as attribute-level replication. Unfortunately, not all Seam users have the good fortune to be working in an environment that supports EJB 3.0. So, for session and conversation scoped `JavaBean` and entity bean components, Seam provides an extra layer of cluster-safe state management over the top of the web container session clustering.

For session or conversation scoped `JavaBean` components, Seam automatically forces replication to occur by calling `setAttribute()` once in every request that the component was invoked by the application. Of course, this strategy is inefficient for read-mostly components. You can control this behavior by implementing the `org.jboss.seam.core.Mutable` interface, or by extending `org.jboss.seam.core.AbstractMutable`, and writing your own dirty-checking logic inside the component. For example,

```
@Name("account")
public class Account extends AbstractMutable
{
    private BigDecimal balance;

    public void setBalance(BigDecimal balance)
    {
        setDirty(this.balance, balance);
        this.balance = balance;
    }

    public BigDecimal getBalance()
    {
        return balance;
    }
}
```

```

    {
        return balance;
    }

    ...
}

```

Or, you can use the `@ReadOnly` annotation to achieve a similar effect:

```

@Name("account")
public class Account
{
    private BigDecimal balance;

    public void setBalance(BigDecimal balance)
    {
        this.balance = balance;
    }

    @ReadOnly
    public BigDecimal getBalance()
    {
        return balance;
    }

    ...
}

```

For session or conversation scoped entity bean components, Seam automatically forces replication to occur by calling `setAttribute()` once in every request, *unless the (conversation-scoped) entity is currently associated with a Seam-managed persistence context, in which case no replication is needed*. This strategy is not necessarily efficient, so session or conversation scope entity beans should be used with care. You can always write a stateful session bean or JavaBean component to "manage" the entity bean instance. For example,

```

@Stateful
@Name("account")
public class AccountManager extends AbstractMutable
{
    private Account account; // an entity bean

    @Unwrap
    public void getAccount()
    {
        return account;
    }

    ...
}

```

Note that the `EntityHome` class in the Seam Application Framework provides a great example of managing an entity bean instance using a Seam component.

3.8. Factory and manager components

We often need to work with objects that are not Seam components. But we still want to be able to inject them into our components using `@In` and use them in value and method binding expressions, etc. Sometimes, we even need to tie them into the Seam context lifecycle (`@Destroy`, for example). So the Seam contexts can contain objects which are not Seam components, and Seam provides a couple of nice features that make it easier to

work with non-component objects bound to contexts.

The *factory component pattern* lets a Seam component act as the instantiator for a non-component object. A *factory method* will be called when a context variable is referenced but has no value bound to it. We define factory methods using the `@Factory` annotation. The factory method binds a value to the context variable, and determines the scope of the bound value. There are two styles of factory method. The first style returns a value, which is bound to the context by Seam:

```
@Factory(scope=CONVERSATION)
public List<Customer> getCustomerList() {
    return ... ;
}
```

The second style is a method of type void which binds the value to the context variable itself:

```
@DataModel List<Customer> customerList;

@Factory("customerList")
public void initCustomerList() {
    customerList = ... ;
}
```

In both cases, the factory method is called when we reference the `customerList` context variable and its value is null, and then has no further part to play in the lifecycle of the value. An even more powerful pattern is the *manager component pattern*. In this case, we have a Seam component that is bound to a context variable, that manages the value of the context variable, while remaining invisible to clients.

A manager component is any component with an `@Unwrap` method. This method returns the value that will be visible to clients, and is called *every time* a context variable is referenced.

```
@Name("customerList")
@Scope(CONVERSATION)
public class CustomerListManager
{
    ...

    @Unwrap
    public List<Customer> getCustomerList() {
        return ... ;
    }
}
```

The manager component pattern is especially useful if we have an object where you need more control over the lifecycle of the component. For example, if you have a heavyweight object that needs a cleanup operation when the context ends you could `@Unwrap` the object, and perform cleanup in the `@Destroy` method of the manager component.

```
@Name("hens")
@Scope(APPLICATION)
public class HenHouse {

    Set<Hen> hens;

    @In(required=false) Hen hen;

    @Unwrap
    public List<Hen> getHens() {
        if (hens == null) {
            // Setup our hens
        }
        return hens;
    }
}
```

```
@Observer({"chickBorn", "chickenBoughtAtMarket"})
public addHen() {
    hens.add(hen);
}

@Observer("chickenSoldAtMarket")
public removeHen() {
    hens.remove(hen);
}

@Observer("foxGetsIn")
public removeAllHens() {
    hens.clear();
}
...
}
```

Here the managed component observes many events which change the underlying object. The component manages these actions itself, and because the object is unwrapped on every access, a consistent view is provided.

Chapter 4. Configuring Seam components

The philosophy of minimizing XML-based configuration is extremely strong in Seam. Nevertheless, there are various reasons why we might want to configure a Seam component using XML: to isolate deployment-specific information from the Java code, to enable the creation of re-usable frameworks, to configure Seam's built-in functionality, etc. Seam provides two basic approaches to configuring components: configuration via property settings in a properties file or in `web.xml`, and configuration via `components.xml`.

4.1. Configuring components via property settings

Seam components may be provided with configuration properties either via servlet context parameters, or via a properties file named `seam.properties` in the root of the classpath.

The configurable Seam component must expose JavaBeans-style property setter methods for the configurable attributes. If a Seam component named `com.jboss.myapp.settings` has a setter method named `setLocale()`, we can provide a property named `com.jboss.myapp.settings.locale` in the `seam.properties` file or as a servlet context parameter, and Seam will set the value of the `locale` attribute whenever it instantiates the component.

The same mechanism is used to configure Seam itself. For example, to set the conversation timeout, we provide a value for `org.jboss.seam.core.manager.conversationTimeout` in `web.xml` or `seam.properties`. (There is a built-in Seam component named `org.jboss.seam.core.manager` with a setter method named `setConversationTimeout()`.)

4.2. Configuring components via `components.xml`

The `components.xml` file is a bit more powerful than property settings. It lets you:

- Configure components that have been installed automatically—including both built-in components, and application components that have been annotated with the `@Name` annotation and picked up by Seam's deployment scanner.
- Install classes with no `@Name` annotation as Seam components—this is most useful for certain kinds of infrastructural components which can be installed multiple times different names (for example Seam-managed persistence contexts).
- Install components that *do* have a `@Name` annotation but are not installed by default because of an `@Install` annotation that indicates the component should not be installed.
- Override the scope of a component.

A `components.xml` file may appear in one of three different places:

- The `WEB-INF` directory of a war.
- The `META-INF` directory of a jar.
- Any directory of a jar that contains classes with an `@Name` annotation.

Usually, Seam components are installed when the deployment scanner discovers a class with a `@Name` annota-

tion sitting in an archive with a `seam.properties` file or a `META-INF/components.xml` file. (Unless the component has an `@Install` annotation indicating it should not be installed by default.) The `components.xml` file lets us handle special cases where we need to override the annotations.

For example, the following `components.xml` file installs jBPM:

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
            xmlns:bpm="http://jboss.com/products/seam/bpm">
  <bpm:jbp/>
</components>
```

This example does the same thing:

```
<components>
  <component class="org.jboss.seam.bpm.Jbpm"/>
</components>
```

This one installs and configures two different Seam-managed persistence contexts:

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:persistence="http://jboss.com/products/seam/persistence">

  <persistence:managed-persistence-context name="customerDatabase"
                                          persistence-unit-jndi-name="java:/customerEntityManagerFactory"/>

  <persistence:managed-persistence-context name="accountingDatabase"
                                          persistence-unit-jndi-name="java:/accountingEntityManagerFactory"/>

</components>
```

As does this one:

```
<components>
  <component name="customerDatabase"
            class="org.jboss.seam.persistence.ManagedPersistenceContext">
    <property name="persistenceUnitJndiName">java:/customerEntityManagerFactory</property>
  </component>

  <component name="accountingDatabase"
            class="org.jboss.seam.persistence.ManagedPersistenceContext">
    <property name="persistenceUnitJndiName">java:/accountingEntityManagerFactory</property>
  </component>
</components>
```

This example creates a session-scoped Seam-managed persistence context (this is not recommended in practice):

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:persistence="http://jboss.com/products/seam/persistence">

  <persistence:managed-persistence-context name="productDatabase"
                                          scope="session"
                                          persistence-unit-jndi-name="java:/productEntityManagerFactory"/>

</components>
```

```
<components>

  <component name="productDatabase"
            scope="session"
            class="org.jboss.seam.persistence.ManagedPersistenceContext">
    <property name="persistenceUnitJndiName">java:/productEntityManagerFactory</property>
  </component>

</components>
```

```
</components>
```

It is common to use the `auto-create` option for infrastructural objects like persistence contexts, which saves you from having to explicitly specify `create=true` when you use the `@In` annotation.

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:persistence="http://jboss.com/products/seam/persistence"

  <persistence:managed-persistence-context name="productDatabase"
    auto-create="true"
    persistence-unit-jndi-name="java:/productEntityManagerFactory"/>

</components>
```

```
<components>

  <component name="productDatabase"
    auto-create="true"
    class="org.jboss.seam.persistence.ManagedPersistenceContext">
    <property name="persistenceUnitJndiName">java:/productEntityManagerFactory</property>
  </component>

</components>
```

The `<factory>` declaration lets you specify a value or method binding expression that will be evaluated to initialize the value of a context variable when it is first referenced.

```
<components>

  <factory name="contact" method="#{contactManager.loadContact}" scope="CONVERSATION"/>

</components>
```

You can create an "alias" (a second name) for a Seam component like so:

```
<components>

  <factory name="user" value="#{actor}" scope="STATELESS"/>

</components>
```

You can even create an "alias" for a commonly used expression:

```
<components>

  <factory name="contact" value="#{contactManager.contact}" scope="STATELESS"/>

</components>
```

It is especially common to see the use of `auto-create="true"` with the `<factory>` declaration:

```
<components>

  <factory name="session" value="#{entityManager.delegate}" scope="STATELESS" auto-create="true"/>

</components>
```

Sometimes we want to reuse the same `components.xml` file with minor changes during both deployment and testing. Seam lets you place wildcards of the form `@wildcard@` in the `components.xml` file which can be replaced either by your Ant build script (at deployment time) or by providing a file named `compon-`

ents.properties in the classpath (at development time). You'll see this approach used in the Seam examples.

4.3. Fine-grained configuration files

If you have a large number of components that need to be configured in XML, it makes much more sense to split up the information in components.xml into many small files. Seam lets you put configuration for a class named, for example, com.helloworld.Hello in a resource named com/helloworld/Hello.component.xml. (You might be familiar with this pattern, since it is the same one we use in Hibernate.) The root element of the file may be either a <components> or <component> element.

The first option lets you define multiple components in the file:

```
<components>
  <component class="com.helloworld.Hello" name="hello">
    <property name="name">#{user.name}</property>
  </component>
  <factory name="message" value="#{hello.message}"/>
</components>
```

The second option only lets you define or configure one component, but is less noisy:

```
<component name="hello">
  <property name="name">#{user.name}</property>
</component>
```

In the second option, the class name is implied by the file in which the component definition appears.

Alternatively, you may put configuration for all classes in the com.helloworld package in com/helloworld/components.xml.

4.4. Configurable property types

Properties of string, primitive or primitive wrapper type may be configured just as you would expect:

```
org.jboss.seam.core.manager.conversationTimeout 60000
```

```
<core:manager conversation-timeout="60000"/>
```

```
<component name="org.jboss.seam.core.manager">
  <property name="conversationTimeout">60000</property>
</component>
```

Arrays, sets and lists of strings or primitives are also supported:

```
org.jboss.seam.bpm.jbpm.processDefinitions order.jpdl.xml, return.jpdl.xml, inventory.jpdl.xml
```

```
<bpm:jbpm>
  <bpm:process-definitions>
    <value>order.jpdl.xml</value>
    <value>return.jpdl.xml</value>
    <value>inventory.jpdl.xml</value>
  </bpm:process-definitions>
</bpm:jbpm>
```

```
<component name="org.jboss.seam.bpm.jbpm">
  <property name="processDefinitions">
```

```

    <value>order.jpdl.xml</value>
    <value>return.jpdl.xml</value>
    <value>inventory.jpdl.xml</value>
  </property>
</component>

```

Even maps with String-valued keys and string or primitive values are supported:

```

<component name="issueEditor">
  <property name="issueStatuses">
    <key>open</key> <value>open issue</value>
    <key>resolved</key> <value>issue resolved by developer</value>
    <key>closed</key> <value>resolution accepted by user</value>
  </property>
</component>

```

Finally, you may wire together components using a value-binding expression. Note that this is quite different to injection using `@In`, since it happens at component instantiation time instead of invocation time. It is therefore much more similar to the dependency injection facilities offered by traditional IoC containers like JSF or Spring.

```

<drools:managed-working-memory name="policyPricingWorkingMemory" rule-base="#{policyPricingRules}"/>

```

```

<component name="policyPricingWorkingMemory"
  class="org.jboss.seam.drools.ManagedWorkingMemory">
  <property name="ruleBase">#{policyPricingRules}</property>
</component>

```

4.5. Using XML Namespaces

Throughout the examples, there have been two competing ways of declaring components: with and without the use of XML namespaces. The following shows a typical `components.xml` file without namespaces:

```

<?xml version="1.0" encoding="UTF-8"?>
<components xmlns="http://jboss.com/products/seam/components"
  xsi:schemaLocation="http://jboss.com/products/seam/components http://jboss.com/products/seam/components-2.0.xsd">

  <component class="org.jboss.seam.core.init">
    <property name="debug">true</property>
    <property name="jndiPattern">@jndiPattern@</property>
  </component>

</components>

```

As you can see, this is somewhat verbose. Even worse, the component and attribute names cannot be validated at development time.

The namespaced version looks like this:

```

<?xml version="1.0" encoding="UTF-8"?>
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://jboss.com/products/seam/core http://jboss.com/products/seam/core-2.0.xsd
    http://jboss.com/products/seam/components http://jboss.com/products/seam/components-2.0.xsd">

  <core:init debug="true" jndi-pattern="@jndiPattern@" />

</components>

```

Even though the schema declarations are verbose, the actual XML content is lean and easy to understand. The schemas provide detailed information about each component and the attributes available, allowing XML editors to offer intelligent autocomplete. The use of namespaced elements makes generating and maintaining correct `components.xml` files much simpler.

Now, this works great for the built-in Seam components, but what about user components? There are two options. First, Seam supports mixing the two models, allowing the use of the generic `<component>` declarations for user components, along with namespaced declarations for built-in components. But even better, Seam allows you to quickly declare namespaces for your own components.

Any Java package can be associated with an XML namespace by annotating the package with the `@Namespace` annotation. (Package-level annotations are declared in a file named `package-info.java` in the package directory.) Here is an example from the `seampay` demo:

```
@Namespace(value="http://jboss.com/products/seam/examples/seampay")
package org.jboss.seam.example.seampay;

import org.jboss.seam.annotations.Namespace;
```

That is all you need to do to use the namespaced style in `components.xml`! Now we can write:

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:pay="http://jboss.com/products/seam/examples/seampay"
            ... >

    <pay:payment-home new-instance="#{newPayment}"
                    created-message="Created a new payment to #{newPayment.payee}" />

    <pay:payment name="newPayment"
                payee="Somebody"
                account="#{selectedAccount}"
                payment-date="#{currentDatetime}"
                created-date="#{currentDatetime}" />
    ...
</components>
```

Or:

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:pay="http://jboss.com/products/seam/examples/seampay"
            ... >

    <pay:payment-home>
        <pay:new-instance>#{newPayment}</pay:new-instance>
        <pay:created-message>Created a new payment to #{newPayment.payee}</pay:created-message>
    </pay:payment-home>

    <pay:payment name="newPayment">
        <pay:payee>Somebody</pay:payee>
        <pay:account>#{selectedAccount}</pay:account>
        <pay:payment-date>#{currentDatetime}</pay:payment-date>
        <pay:created-date>#{currentDatetime}</pay:created-date>
    </pay:payment>
    ...
</components>
```

These examples illustrate the two usage models of a namespaced element. In the first declaration, the `<pay:payment-home>` references the `paymentHome` component:

```
package org.jboss.seam.example.seampay;
...
@Name("paymentHome")
public class PaymentController
```

```
    extends EntityHome<Payment>
{
    ...
}
```

The element name is the hyphenated form of the component name. The attributes of the element are the hyphenated form of the property names.

In the second declaration, the `<pay:payment>` element refers to the `Payment` class in the `org.jboss.seam.example.seampay` package. In this case `Payment` is an entity that is being declared as a Seam component:

```
package org.jboss.seam.example.seampay;
...
@Entity
public class Payment
    implements Serializable
{
    ...
}
```

If we want validation and autocompletion to work for user-defined components, we will need a schema. Seam does not yet provide a mechanism to automatically generate a schema for a set of components, so it is necessary to generate one manually. The schema definitions for the standard Seam packages can be used for guidance.

The following are the namespaces used by Seam:

- `components` — <http://jboss.com/products/seam/components>
- `core` — <http://jboss.com/products/seam/core>
- `drools` — <http://jboss.com/products/seam/drools>
- `framework` — <http://jboss.com/products/seam/framework>
- `jms` — <http://jboss.com/products/seam/jms>
- `remoting` — <http://jboss.com/products/seam/remoting>
- `theme` — <http://jboss.com/products/seam/theme>
- `security` — <http://jboss.com/products/seam/security>
- `mail` — <http://jboss.com/products/seam/mail>
- `web` — <http://jboss.com/products/seam/web>
- `pdf` — <http://jboss.com/products/seam/pdf>
- `spring` — <http://jboss.com/products/seam/spring>

Chapter 5. Events, interceptors and exception handling

Complementing the contextual component model, there are two further basic concepts that facilitate the extreme loose-coupling that is the distinctive feature of Seam applications. The first is a strong event model where events may be mapped to event listeners via JSF-like method binding expressions. The second is the pervasive use of annotations and interceptors to apply cross-cutting concerns to components which implement business logic.

5.1. Seam events

The Seam component model was developed for use with *event-driven applications*, specifically to enable the development of fine-grained, loosely-coupled components in a fine-grained eventing model. Events in Seam come in several types, most of which we have already seen:

- JSF events
- jBPM transition events
- Seam page actions
- Seam component-driven events
- Seam contextual events

All of these various kinds of events are mapped to Seam components via JSF EL method binding expressions. For a JSF event, this is defined in the JSF template:

```
<h:commandButton value="Click me!" action="#{helloWorld.sayHello}"/>
```

For a jBPM transition event, it is specified in the jBPM process definition or pageflow definition:

```
<start-page name="hello" view-id="/hello.jsp">
  <transition to="hello">
    <action expression="#{helloWorld.sayHello}"/>
  </transition>
</start-page>
```

You can find out more information about JSF events and jBPM events elsewhere. Let's concentrate for now upon the two additional kinds of events defined by Seam.

5.1.1. Page actions

A Seam page action is an event that occurs just before we render a page. We declare page actions in `WEB-INF/pages.xml`. We can define a page action for either a particular JSF view id:

```
<pages>
  <page view-id="/hello.jsp" action="#{helloWorld.sayHello}"/>
</pages>
```

Or we can use a `*` wildcard as a suffix to the `view-id` to specify an action that applies to all view ids that match

the pattern:

```
<pages>
  <page view-id="/hello/*" action="#{helloWorld.sayHello}"/>
</pages>
```

If multiple wildcarded page actions match the current view-id, Seam will call all the actions, in order of least-specific to most-specific.

The page action method can return a JSF outcome. If the outcome is non-null, Seam will use the defined navigation rules to navigate to a view.

Furthermore, the view id mentioned in the `<page>` element need not correspond to a real JSP or Facelets page! So, we can reproduce the functionality of a traditional action-oriented framework like Struts or WebWork using page actions. For example:

```
TODO: translate struts action into page action
```

This is quite useful if you want to do complex things in response to non-faces requests (for example, HTTP GET requests).

Multiple or conditional page actions may be specified using the `<action>` tag:

```
<pages>
  <page view-id="/hello.jsp">
    <action execute="#{helloWorld.sayHello}" if="#{not validation.failed}"/>
    <action execute="#{hitCount.increment}"/>
  </page>
</pages>
```

Page parameters

A JSF faces request (a form submission) encapsulates both an "action" (a method binding) and "parameters" (input value bindings). A page action might also need parameters!

Since GET requests are bookmarkable, page parameters are passed as human-readable request parameters. (Unlike JSF form inputs, which are anything but!)

You can use page parameters with or without an action method.

Mapping request parameters to the model

Seam lets us provide a value binding that maps a named request parameter to an attribute of a model object.

```
<pages>
  <page view-id="/hello.jsp" action="#{helloWorld.sayHello}">
    <param name="firstName" value="#{person.firstName}"/>
    <param name="lastName" value="#{person.lastName}"/>
  </page>
</pages>
```

The `<param>` declaration is bidirectional, just like a value binding for a JSF input:

- When a non-faces (GET) request for the view id occurs, Seam sets the value of the named request parameter onto the model object, after performing appropriate type conversions.
- Any `<s:link>` or `<s:button>` transparently includes the request parameter. The value of the parameter is

determined by evaluating the value binding during the render phase (when the `<s:link>` is rendered).

- Any navigation rule with a `<redirect/>` to the view id transparently includes the request parameter. The value of the parameter is determined by evaluating the value binding at the end of the invoke application phase.
- The value is transparently propagated with any JSF form submission for the page with the given view id. This means that view parameters behave like PAGE-scoped context variables for faces requests.

The essential idea behind all this is that *however* we get from any other page to `/hello.jsp` (or from `/hello.jsp` back to `/hello.jsp`), the value of the model attribute referred to in the value binding is "remembered", without the need for a conversation (or other server-side state).

Propagating request parameters

If just the name attribute is specified then the request parameter is propagated using the PAGE context (it isn't mapped to model property).

```
<pages>
  <page view-id="/hello.jsp" action="#{helloWorld.sayHello}">
    <param name="firstName" />
    <param name="lastName" />
  </page>
</pages>
```

Propagation of page parameters is especially useful if you want to build multi-layer master-detail CRUD pages. You can use it to "remember" which view you were previously on (e.g. when pressing the Save button), and which entity you were editing.

- Any `<s:link>` or `<s:button>` transparently propagates the request parameter if that parameter is listed as a page parameter for the view.
- The value is transparently propagated with any JSF form submission for the page with the given view id. (This means that view parameters behave like PAGE-scoped context variables for faces requests.)

This all sounds pretty complex, and you're probably wondering if such an exotic construct is really worth the effort. Actually, the idea is very natural once you "get it". It is definitely worth taking the time to understand this stuff. Page parameters are the most elegant way to propagate state across a non-faces request. They are especially cool for problems like search screens with bookmarkable results pages, where we would like to be able to write our application code to handle both POST and GET requests with the same code. Page parameters eliminate repetitive listing of request parameters in the view definition and make redirects much easier to code.

Conversion and Validation

You can specify a JSF converter for complex model properties:

```
<pages>
  <page view-id="/calculator.jsp" action="#{calculator.calculate}">
    <param name="x" value="#{calculator.lhs}" />
    <param name="y" value="#{calculator.rhs}" />
    <param name="op" converterId="com.my.calculator.OperatorConverter" value="#{calculator.op}" />
  </page>
</pages>
```

Alternatively:

```
<pages>
```

```
<page view-id="/calculator.jsp" action="#{calculator.calculate}">
  <param name="x" value="#{calculator.lhs}"/>
  <param name="y" value="#{calculator.rhs}"/>
  <param name="op" converter="#{operatorConverter}" value="#{calculator.op}"/>
</page>
</pages>
```

JSF validators, and `required="true"` may also be used:

```
<pages>
  <page view-id="/blog.xhtml">
    <param name="date"
      value="#{blog.date}"
      validatorId="com.my.blog.PastDate"
      required="true"/>
  </page>
</pages>
```

Alternatively:

```
<pages>
  <page view-id="/blog.xhtml">
    <param name="date"
      value="#{blog.date}"
      validator="#{pastDateValidator}"
      required="true"/>
  </page>
</pages>
```

Even better, model-based Hibernate validator annotations are automatically recognized and validated.

When type conversion or validation fails, a global `FacesMessage` is added to the `FacesContext`.

Navigation

You can use standard JSF navigation rules defined in `faces-config.xml` in a Seam application. However, JSF navigation rules have a number of annoying limitations:

- It is not possible to specify request parameters to be used when redirecting.
- It is not possible to begin or end conversations from a rule.
- Rules work by evaluating the return value of the action method; it is not possible to evaluate an arbitrary EL expression.

A further problem is that "orchestration" logic gets scattered between `pages.xml` and `faces-config.xml`. It's better to unify this logic into `pages.xml`.

This JSF navigation rule:

```
<navigation-rule>
  <from-view-id>/editDocument.xhtml</from-view-id>

  <navigation-case>
    <from-action>#{documentEditor.update}</from-action>
    <from-outcome>success</from-outcome>
    <to-view-id>/viewDocument.xhtml</to-view-id>
    <redirect/>
  </navigation-case>
</navigation-rule>
```

Can be rewritten as follows:

```
<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <rule if-outcome="success">
            <redirect view-id="/viewDocument.xhtml"/>
        </rule>
    </navigation>

</page>
```

But it would be even nicer if we didn't have to pollute our `DocumentEditor` component with string-valued return values (the JSF outcomes). So Seam lets us write:

```
<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}"
                evaluate="#{documentEditor.errors.size}">
        <rule if-outcome="0">
            <redirect view-id="/viewDocument.xhtml"/>
        </rule>
    </navigation>

</page>
```

Or even:

```
<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <rule if="#{documentEditor.errors.empty}">
            <redirect view-id="/viewDocument.xhtml"/>
        </rule>
    </navigation>

</page>
```

The first form evaluates a value binding to determine the outcome value to be used by the subsequent rules. The second approach ignores the outcome and evaluates a value binding for each possible rule.

Of course, when an update succeeds, we probably want to end the current conversation. We can do that like this:

```
<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <rule if="#{documentEditor.errors.empty}">
            <end-conversation/>
            <redirect view-id="/viewDocument.xhtml"/>
        </rule>
    </navigation>

</page>
```

As we've ended conversation any subsequent requests won't know which document we are interested in. We can pass the document id as a request parameter which also makes the view bookmarkable:

```
<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <rule if="#{documentEditor.errors.empty}">
            <end-conversation/>
            <redirect view-id="/viewDocument.xhtml">
```

```

        <param name="documentId" value="#{documentEditor.documentId}"/>
    </redirect>
</rule>
</navigation>
</page>

```

Null outcomes are a special case in JSF. The null outcome is interpreted to mean "redisplay the page". The following navigation rule matches any non-null outcome, but *not* the null outcome:

```

<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <rule>
            <render view-id="/viewDocument.xhtml"/>
        </rule>
    </navigation>

</page>

```

If you want to perform navigation when a null outcome occurs, use the following form instead:

```

<page view-id="/editDocument.xhtml">

    <navigation from-action="#{documentEditor.update}">
        <render view-id="/viewDocument.xhtml"/>
    </navigation>

</page>

```

The view-id may be given as a JSF EL expression:

```

<page view-id="/editDocument.xhtml">

    <navigation>
        <rule if-outcome="success">
            <redirect view-id="#{userAgent}/displayDocument.xhtml"/>
        </rule>
    </navigation>

</page>

```

Fine-grained files for definition of navigation, page actions and parameters

If you have a lot of different page actions and page parameters, or even just a lot of navigation rules, you will almost certainly want to split the declarations up over multiple files. You can define actions and parameters for a page with the view id `/calc/calculator.jsp` in a resource named `calc/calculator.page.xml`. The root element in this case is the `<page>` element, and the view id is implied:

```

<page action="#{calculator.calculate}">
    <param name="x" value="#{calculator.lhs}"/>
    <param name="y" value="#{calculator.rhs}"/>
    <param name="op" converter="#{operatorConverter}" value="#{calculator.op}"/>
</page>

```

5.1.2. Component-driven events

Seam components can interact by simply calling each others methods. Stateful components may even implement the observer/observable pattern. But to enable components to interact in a more loosely-coupled fashion than is possible when the components call each others methods directly, Seam provides *component-driven*

events.

We specify event listeners (observers) in `components.xml`.

```
<components>
  <event type="hello">
    <action execute="#{helloListener.sayHelloBack}"/>
    <action execute="#{logger.logHello}"/>
  </event>
</components>
```

Where the *event type* is just an arbitrary string.

When an event occurs, the actions registered for that event will be called in the order they appear in `components.xml`. How does a component raise an event? Seam provides a built-in component for this.

```
@Name("helloWorld")
public class HelloWorld {
    public void sayHello() {
        FacesMessages.instance().add("Hello World!");
        Events.instance().raiseEvent("hello");
    }
}
```

Or you can use an annotation.

```
@Name("helloWorld")
public class HelloWorld {
    @RaiseEvent("hello")
    public void sayHello() {
        FacesMessages.instance().add("Hello World!");
    }
}
```

Notice that this event producer has no dependency upon event consumers. The event listener may now be implemented with absolutely no dependency upon the producer:

```
@Name("helloListener")
public class HelloListener {
    public void sayHelloBack() {
        FacesMessages.instance().add("Hello to you too!");
    }
}
```

The method binding defined in `components.xml` above takes care of mapping the event to the consumer. If you don't like futzing about in the `components.xml` file, you can use an annotation instead:

```
@Name("helloListener")
public class HelloListener {
    @Observer("hello")
    public void sayHelloBack() {
        FacesMessages.instance().add("Hello to you too!");
    }
}
```

You might wonder why I've not mentioned anything about event objects in this discussion. In Seam, there is no need for an event object to propagate state between event producer and listener. State is held in the Seam contexts, and is shared between components. However, if you really want to pass an event object, you can:

```
@Name("helloWorld")
public class HelloWorld {
    private String name;
    public void sayHello() {
```

```
        FacesMessages.instance().add("Hello World, my name is #0.", name);
        Events.instance().raiseEvent("hello", name);
    }
}
```

```
@Name("helloListener")
public class HelloListener {
    @Observer("hello")
    public void sayHelloBack(String name) {
        FacesMessages.instance().add("Hello #0!", name);
    }
}
```

5.1.3. Contextual events

Seam defines a number of built-in events that the application can use to perform special kinds of framework integration. The events are:

- `org.jboss.seam.validationFailed` — called when JSF validation fails
- `org.jboss.seam.noConversation` — called when there is no long running conversation and a long running conversation is required
- `org.jboss.seam.preSetVariable.<name>` — called when the context variable `<name>` is set
- `org.jboss.seam.postSetVariable.<name>` — called when the context variable `<name>` is set
- `org.jboss.seam.preRemoveVariable.<name>` — called when the context variable `<name>` is unset
- `org.jboss.seam.postRemoveVariable.<name>` — called when the context variable `<name>` is unset
- `org.jboss.seam.preDestroyContext.<SCOPE>` — called before the `<SCOPE>` context is destroyed
- `org.jboss.seam.postDestroyContext.<SCOPE>` — called after the `<SCOPE>` context is destroyed
- `org.jboss.seam.beginConversation` — called whenever a long-running conversation begins
- `org.jboss.seam.endConversation` — called whenever a long-running conversation ends
- `org.jboss.seam.beginPageflow.<name>` — called when the pageflow `<name>` begins
- `org.jboss.seam.endPageflow.<name>` — called when the pageflow `<name>` ends
- `org.jboss.seam.createProcess.<name>` — called when the process `<name>` is created
- `org.jboss.seam.endProcess.<name>` — called when the process `<name>` ends
- `org.jboss.seam.initProcess.<name>` — called when the process `<name>` is associated with the conversation
- `org.jboss.seam.initTask.<name>` — called when the task `<name>` is associated with the conversation
- `org.jboss.seam.startTask.<name>` — called when the task `<name>` is started
- `org.jboss.seam.endTask.<name>` — called when the task `<name>` is ended
- `org.jboss.seam.postCreate.<name>` — called when the component `<name>` is created

- `org.jboss.seam.preDestroy.<name>` — called when the component `<name>` is destroyed
- `org.jboss.seam.beforePhase` — called before the start of a JSF phase
- `org.jboss.seam.afterPhase` — called after the end of a JSF phase
- `org.jboss.seam.postInitialization` — called when Seam has initialized and started up all components
- `org.jboss.seam.postAuthenticate.<name>` — called after a user is authenticated
- `org.jboss.seam.preAuthenticate.<name>` — called before attempting to authenticate a user
- `org.jboss.seam.notLoggedIn` — called there is no authenticated user and authentication is required
- `org.jboss.seam.rememberMe` — occurs when Seam security detects the username in a cookie
- `org.jboss.seam.exceptionHandled.<type>` — called when an uncaught exception is handled by Seam
- `org.jboss.seam.exceptionHandled` — called when an uncaught exception is handled by Seam
- `org.jboss.seam.exceptionNotHandled` — called when there was no handler for an uncaught exception
- `org.jboss.seam.afterTransactionSuccess` — called when a transaction succeeds in the Seam Application Framework
- `org.jboss.seam.afterTransactionSuccess.<name>` — called when a transaction succeeds in the Seam Application Framework which manages an entity called `<name>`

Seam components may observe any of these events in just the same way they observe any other component-driven events.

5.2. Seam interceptors

EJB 3.0 introduced a standard interceptor model for session bean components. To add an interceptor to a bean, you need to write a class with a method annotated `@AroundInvoke` and annotate the bean with an `@Interceptors` annotation that specifies the name of the interceptor class. For example, the following interceptor checks that the user is logged in before allowing invoking an action listener method:

```
public class LoggedInInterceptor {

    @AroundInvoke
    public Object checkLoggedIn(InvocationContext invocation) throws Exception {

        boolean isLoggedIn = Contexts.getSessionContext().get("loggedIn")!=null;
        if (isLoggedIn) {
            //the user is already logged in
            return invocation.proceed();
        }
        else {
            //the user is not logged in, fwd to login page
            return "login";
        }
    }

}
```

To apply this interceptor to a session bean which acts as an action listener, we must annotate the session bean `@Interceptors(LoggedInInterceptor.class)`. This is a somewhat ugly annotation. Seam builds upon the in-

interceptor framework in EJB3 by allowing you to use `@Interceptors` as a meta-annotation for class level interceptors (those annotated `@Target(TYPE)`). In our example, we would create an `@LoggedIn` annotation, as follows:

```
@Target(TYPE)
@Retention(RUNTIME)
@Interceptors(LoginInterceptor.class)
public @interface LoggedIn {}
```

We can now simply annotate our action listener bean with `@LoggedIn` to apply the interceptor.

```
@Stateless
@Name("changePasswordAction")
@LoggedIn
@Interceptors(SeamInterceptor.class)
public class ChangePasswordAction implements ChangePassword {

    ...

    public String changePassword() { ... }

}
```

If interceptor ordering is important (it usually is), you can add `@Interceptor` annotations to your interceptor classes to specify a partial order of interceptors.

```
@Interceptor(around={BijectionInterceptor.class,
                    ValidationInterceptor.class,
                    ConversationInterceptor.class},
            within=RemoveInterceptor.class)
public class LoggedInInterceptor
{
    ...
}
```

You can even have a "client-side" interceptor, that runs around any of the built-in functionality of EJB3:

```
@Interceptor(type=CLIENT)
public class LoggedInInterceptor
{
    ...
}
```

EJB interceptors are stateful, with a lifecycle that is the same as the component they intercept. For interceptors which do not need to maintain state, Seam lets you get a performance optimization by specifying `@Interceptor(stateless=true)`.

Much of the functionality of Seam is implemented as a set of built-in Seam interceptors, including the interceptors named in the previous example. You don't have to explicitly specify these interceptors by annotating your components; they exist for all interceptable Seam components.

You can even use Seam interceptors with JavaBean components, not just EJB3 beans!

EJB defines interception not only for business methods (using `@AroundInvoke`), but also for the lifecycle methods `@PostConstruct`, `@PreDestroy`, `@PrePassivate` and `@PostActive`. Seam supports all these lifecycle methods on both component and interceptor not only for EJB3 beans, but also for JavaBean components (except `@PreDestroy` which is not meaningful for JavaBean components).

5.3. Managing exceptions

JSF is surprisingly limited when it comes to exception handling. As a partial workaround for this problem, Seam lets you define how a particular class of exception is to be treated by annotating the exception class, or declaring the exception class in an XML file. This facility is meant to be combined with the EJB 3.0-standard `@ApplicationException` annotation which specifies whether the exception should cause a transaction rollback.

5.3.1. Exceptions and transactions

EJB specifies well-defined rules that let us control whether an exception immediately marks the current transaction for rollback when it is thrown by a business method of the bean: *system exceptions* always cause a transaction rollback, *application exceptions* do not cause a rollback by default, but they do if `@ApplicationException(rollback=true)` is specified. (An application exception is any checked exception, or any unchecked exception annotated `@ApplicationException`. A system exception is any unchecked exception without an `@ApplicationException` annotation.)

Note that there is a difference between marking a transaction for rollback, and actually rolling it back. The exception rules say that the transaction should be marked rollback only, but it may still be active after the exception is thrown.

Seam applies the EJB 3.0 exception rollback rules also to Seam JavaBean components.

But these rules only apply in the Seam component layer. What about an exception that is uncaught and propagates out of the Seam component layer, and out of the JSF layer? Well, it is always wrong to leave a dangling transaction open, so Seam rolls back any active transaction when an exception occurs and is uncaught in the Seam component layer.

5.3.2. Enabling Seam exception handling

To enable Seam's exception handling, we need to make sure we have the master servlet filter declared in `web.xml`:

```
<filter>
  <filter-name>Seam Filter</filter-name>
  <filter-class>org.jboss.seam.servlet.SeamFilter</filter-class>
</filter>

<filter-mapping>
  <filter-name>Seam Filter</filter-name>
  <url-pattern>*.seam</url-pattern>
</filter-mapping>
```

You may also need to disable Facelets development mode in `web.xml` and Seam debug mode in `components.xml` if you want your exception handlers to fire.

5.3.3. Using annotations for exception handling

The following exception results in a HTTP 404 error whenever it propagates out of the Seam component layer. It does not roll back the current transaction immediately when thrown, but the transaction will be rolled back if it the exception is not caught by another Seam component.

```
@HttpError(errorCode=404)
public class ApplicationException extends Exception { ... }
```

This exception results in a browser redirect whenever it propagates out of the Seam component layer. It also ends the current conversation. It causes an immediate rollback of the current transaction.

```
@Redirect(viewId="/failure.xhtml", end=true)
@ApplicationException(rollback=true)
public class UnrecoverableApplicationException extends RuntimeException { ... }
```

Note that `@Redirect` does not work for exceptions which occur during the render phase of the JSF lifecycle.

You can also use EL to specify the `viewId` to redirect to.

This exception results in a redirect, along with a message to the user, when it propagates out of the Seam component layer. It also immediately rolls back the current transaction.

```
@Redirect(viewId="/error.xhtml", message="Unexpected error")
public class SystemException extends RuntimeException { ... }
```

5.3.4. Using XML for exception handling

Since we can't add annotations to all the exception classes we are interested in, Seam also lets us specify this functionality in `pages.xml`.

```
<pages>

  <exception class="javax.persistence.EntityNotFoundException">
    <http-error error-code="404"/>
  </exception>

  <exception class="javax.persistence.PersistenceException">
    <end-conversation/>
    <redirect view-id="/error.xhtml">
      <message>Database access failed</message>
    </redirect>
  </exception>

  <exception>
    <end-conversation/>
    <redirect view-id="/error.xhtml">
      <message>Unexpected failure</message>
    </redirect>
  </exception>

</pages>
```

The last `<exception>` declaration does not specify a class, and is a catch-all for any exception for which handling is not otherwise specified via annotations or in `pages.xml`.

You can also use EL to specify the `view-id` to redirect to.

You can also access the handled exception instance through EL, Seam places it in the conversation context, e.g. to access the message of the exception:

```
...
throw new AuthorizationException("You are not allowed to do this!");

<pages>

  <exception class="org.jboss.seam.security.AuthorizationException">
    <end-conversation/>
    <redirect view-id="/error.xhtml">
      <message severity="WARN">#{org.jboss.seam.handledException.message}</message>
    </redirect>
  </exception>

</pages>
```

`org.jboss.seam.handledException` holds the nested exception that was actually handled by an exception handler. The outermost (wrapper) exception is also available, as `org.jboss.seam.exception`.

5.3.5. Some common exceptions

If you are using JPA:

```
<exception class="javax.persistence.EntityNotFoundException">
  <redirect view-id="/error.xhtml">
    <message>Not found</message>
  </redirect>
</exception>

<exception class="javax.persistence.OptimisticLockException">
  <end-conversation/>
  <redirect view-id="/error.xhtml">
    <message>Another user changed the same data, please try again</message>
  </redirect>
</exception>
```

If you are using the Seam Application Framework:

```
<exception class="org.jboss.seam.framework.EntityNotFoundException">
  <redirect view-id="/error.xhtml">
    <message>Not found</message>
  </redirect>
</exception>
```

If you are using Seam Security:

```
<exception class="org.jboss.seam.security.AuthorizationException">
  <redirect>
    <message>You don't have permission to do this</message>
  </redirect>
</exception>

<exception class="org.jboss.seam.security.NotLoggedInException">
  <redirect view-id="/login.xhtml">
    <message>Please log in first</message>
  </redirect>
</exception>
```

And, for JSF:

```
<exception class="javax.faces.application.ViewExpiredException">
  <redirect view-id="/error.xhtml">
    <message>Your session has timed out, please try again</message>
  </redirect>
</exception>
```

A `ViewExpiredException` occurs if the user posts back to a page once their session has expired. `no-conversation-view-id` and `conversation-required` give you finer grained control over session expiration if you are inside a conversation.

Chapter 6. Conversations and workspace management

It's time to understand Seam's conversation model in more detail.

Historically, the notion of a Seam "conversation" came about as a merger of three different ideas:

- The idea of a *workspace*, which I encountered in a project for the Victorian government in 2002. In this project I was forced to implement workspace management on top of Struts, an experience I pray never to repeat.
- The idea of an *application transaction* with optimistic semantics, and the realization that existing frameworks based around a stateless architecture could not provide effective management of extended persistence contexts. (The Hibernate team is truly fed up with copping the blame for `LazyInitializationExceptions`, which are not really Hibernate's fault, but rather the fault of the extremely limiting persistence context model supported by stateless architectures such as the Spring framework or the traditional *stateless session facade* (anti)pattern in J2EE.)
- The idea of a workflow *task*.

By unifying these ideas and providing deep support in the framework, we have a powerful construct that lets us build richer and more efficient applications with less code than before.

6.1. Seam's conversation model

The examples we have seen so far make use of a very simple conversation model that follows these rules:

- There is always a conversation context active during the apply request values, process validations, update model values, invoke application and render response phases of the JSF request lifecycle.
- At the end of the restore view phase of the JSF request lifecycle, Seam attempts to restore any previous long-running conversation context. If none exists, Seam creates a new temporary conversation context.
- When an `@Begin` method is encountered, the temporary conversation context is promoted to a long running conversation.
- When an `@End` method is encountered, any long-running conversation context is demoted to a temporary conversation.
- At the end of the render response phase of the JSF request lifecycle, Seam stores the contents of a long running conversation context or destroys the contents of a temporary conversation context.
- Any faces request (a JSF postback) will propagate the conversation context. By default, non-faces requests (GET requests, for example) do not propagate the conversation context, but see below for more information on this.
- If the JSF request lifecycle is foreshortened by a redirect, Seam transparently stores and restores the current conversation context—unless the conversation was already ended via `@End(beforeRedirect=true)`.

Seam transparently propagates the conversation context (including the temporary conversation context) across JSF postbacks and redirects. If you don't do anything special, a *non-faces request* (a GET request for example)

will not propagate the conversation context and will be processed in a new temporary conversation. This is usually - but not always - the desired behavior.

If you want to propagate a Seam conversation across a non-faces request, you need to explicitly code the Seam *conversation id* as a request parameter:

```
<a href="main.jsf?conversationId=#{conversation.id}">Continue</a>
```

Or, the more JSF-ish:

```
<h:outputLink value="main.jsf">
  <f:param name="conversationId" value="#{conversation.id}"/>
  <h:outputText value="Continue"/>
</h:outputLink>
```

If you use the Seam tag library, this is equivalent:

```
<h:outputLink value="main.jsf">
  <s:conversationId/>
  <h:outputText value="Continue"/>
</h:outputLink>
```

If you wish to disable propagation of the conversation context for a postback, a similar trick is used:

```
<h:commandLink action="main" value="Exit">
  <f:param name="conversationPropagation" value="none"/>
</h:commandLink>
```

If you use the Seam tag library, this is equivalent:

```
<h:commandLink action="main" value="Exit">
  <s:conversationPropagation type="none"/>
</h:commandLink>
```

Note that disabling conversation context propagation is absolutely not the same thing as ending the conversation.

The `conversationPropagation` request parameter, or the `<s:conversationPropagation>` tag may even be used to begin and end conversation, or begin a nested conversation.

```
<h:commandLink action="main" value="Exit">
  <s:conversationPropagation type="end"/>
</h:commandLink>
```

```
<h:commandLink action="main" value="Select Child">
  <s:conversationPropagation type="nested"/>
</h:commandLink>
```

```
<h:commandLink action="main" value="Select Hotel">
  <s:conversationPropagation type="begin"/>
</h:commandLink>
```

```
<h:commandLink action="main" value="Select Hotel">
  <s:conversationPropagation type="join"/>
</h:commandLink>
```

This conversation model makes it easy to build applications which behave correctly with respect to multi-window operation. For many applications, this is all that is needed. Some complex applications have either or both of the following additional requirements:

- A conversation spans many smaller units of user interaction, which execute serially or even concurrently. The smaller *nested conversations* have their own isolated set of conversation state, and also have access to the state of the outer conversation.
- The user is able to switch between many conversations within the same browser window. This feature is called *workspace management*.

6.2. Nested conversations

A nested conversation is created by invoking a method marked `@Begin(nested=true)` inside the scope of an existing conversation. A nested conversation has its own conversation context, and also has read-only access to the context of the outer conversation. (It can read the outer conversation's context variables, but not write to them.) When an `@End` is subsequently encountered, the nested conversation will be destroyed, and the outer conversation will resume, by "popping" the conversation stack. Conversations may be nested to any arbitrary depth.

Certain user activity (workspace management, or the back button) can cause the outer conversation to be resumed before the inner conversation is ended. In this case it is possible to have multiple concurrent nested conversations belonging to the same outer conversation. If the outer conversation ends before a nested conversation ends, Seam destroys all nested conversation contexts along with the outer context.

A conversation may be thought of as a *continuable state*. Nested conversations allow the application to capture a consistent continuable state at various points in a user interaction, thus insuring truly correct behavior in the face of backbuttoning and workspace management.

TODO: an example to show how a nested conversation prevents bad stuff happening when you backbutton.

Usually, if a component exists in a parent conversation of the current nested conversation, the nested conversation will use the same instance. Occasionally, it is useful to have a different instance in each nested conversation, so that the component instance that exists in the parent conversation is invisible to its child conversations. You can achieve this behavior by annotating the component `@PerNestedConversation`.

6.3. Starting conversations with GET requests

JSF does not define any kind of action listener that is triggered when a page is accessed via a non-faces request (for example, a HTTP GET request). This can occur if the user bookmarks the page, or if we navigate to the page via an `<h:outputLink>`.

Sometimes we want to begin a conversation immediately the page is accessed. Since there is no JSF action method, we can't solve the problem in the usual way, by annotating the action with `@Begin`.

A further problem arises if the page needs some state to be fetched into a context variable. We've already seen two ways to solve this problem. If that state is held in a Seam component, we can fetch the state in a `@Create` method. If not, we can define a `@Factory` method for the context variable.

If none of these options works for you, Seam lets you define a *page action* in the `pages.xml` file.

```
<pages>
  <page view-id="/messageList.jsp" action="#{messageManager.list}"/>
  ...
</pages>
```

This action method is called at the beginning of the render response phase, any time the page is about to be rendered. If a page action returns a non-null outcome, Seam will process any appropriate JSF and Seam navigation rules, possibly resulting in a completely different page being rendered.

If *all* you want to do before rendering the page is begin a conversation, you could use a built-in action method that does just that:

```
<pages>
  <page view-id="/messageList.jsp" action="#{conversation.begin}"/>
  ...
</pages>
```

Note that you can also call this built-in action from a JSF control, and, similarly, you can use `#{conversation.end}` to end conversations.

If you want more control, to join existing conversations or begin a nested conversion, to begin a pageflow or an atomic conversation, you should use the `<begin-conversation>` element.

```
<pages>
  <page view-id="/messageList.jsp">
    <begin-conversation nested="true" pageflow="AddItem"/>
  </page>
  ...
</pages>
```

There is also an `<end-conversation>` element.

```
<pages>
  <page view-id="/home.jsp">
    <end-conversation/>
  </page>
  ...
</pages>
```

To solve the first problem, we now have five options:

- Annotate the `@Create` method with `@Begin`
- Annotate the `@Factory` method with `@Begin`
- Annotate the Seam page action method with `@Begin`
- Use `<begin-conversation>` in `pages.xml`.
- Use `#{conversation.begin}` as the Seam page action method

6.4. Using `<s:link>` and `<s:button>`

JSF command links always perform a form submission via JavaScript, which breaks the web browser's "open in new window" or "open in new tab" feature. In plain JSF, you need to use an `<h:outputLink>` if you need this functionality. But there are two major limitations to `<h:outputLink>`.

- JSF provides no way to attach an action listener to an `<h:outputLink>`.
- JSF does not propagate the selected row of a `DataModel` since there is no actual form submission.

Seam provides the notion of a *page action* to help solve the first problem, but this does nothing to help us with the second problem. We *could* work around this by using the RESTful approach of passing a request parameter and requering for the selected object on the server side. In some cases—such as the Seam blog example application—this is indeed the best approach. The RESTful style supports bookmarking, since it does not require server-side state. In other cases, where we don't care about bookmarks, the use of `@DataModel` and `@DataModelSelection` is just so convenient and transparent!

To fill in this missing functionality, and to make conversation propagation even simpler to manage, Seam provides the `<s:link>` JSF tag.

The link may specify just the JSF view id:

```
<s:link view="/login.xhtml" value="Login"/>
```

Or, it may specify an action method (in which case the action outcome determines the page that results):

```
<s:link action="#{login.logout}" value="Logout"/>
```

If you specify *both* a JSF view id and an action method, the 'view' will be used *unless* the action method returns a non-null outcome:

```
<s:link view="/loggedOut.xhtml" action="#{login.logout}" value="Logout"/>
```

The link automatically propagates the selected row of a `DataModel` using inside `<h:dataTable>`:

```
<s:link view="/hotel.xhtml" action="#{hotelSearch.selectHotel}" value="#{hotel.name}"/>
```

You can leave the scope of an existing conversation:

```
<s:link view="/main.xhtml" propagation="none"/>
```

You can begin, end, or nest conversations:

```
<s:link action="#{issueEditor.viewComment}" propagation="nest"/>
```

If the link begins a conversation, you can even specify a pageflow to be used:

```
<s:link action="#{documentEditor.getDocument}" propagation="begin"
pageflow="EditDocument"/>
```

The `taskInstance` attribute is for use in jBPM task lists:

```
<s:link action="#{documentApproval.approveOrReject}" taskInstance="#{task}"/>
```

(See the DVD Store demo application for examples of this.)

Finally, if you need the "link" to be rendered as a button, use `<s:button>`:

```
<s:button action="#{login.logout}" value="Logout"/>
```

6.5. Success messages

It is quite common to display a message to the user indicating success or failure of an action. It is convenient to

use a JSF `FacesMessage` for this. Unfortunately, a successful action often requires a browser redirect, and JSF does not propagate faces messages across redirects. This makes it quite difficult to display success messages in plain JSF.

The built in conversation-scoped Seam component named `facesMessages` solves this problem. (You must have the Seam redirect filter installed.)

```
@Name("editDocumentAction")
@Stateless
public class EditDocumentBean implements EditDocument {
    @In EntityManager em;
    @In Document document;
    @In FacesMessages facesMessages;

    public String update() {
        em.merge(document);
        facesMessages.add("Document updated");
    }
}
```

Any message added to `facesMessages` is used in the very next render response phase for the current conversation. This even works when there is no long-running conversation since Seam preserves even temporary conversation contexts across redirects.

You can even include JSF EL expressions in a faces message summary:

```
facesMessages.add("Document #{document.title} was updated");
```

You may display the messages in the usual way, for example:

```
<h:messages globalOnly="true"/>
```

6.6. Using an "explicit" conversation id

Ordinarily, Seam generates a meaningless unique id for each conversation in each session. You can customize the id value when you begin the conversation.

This feature can be used to customize the conversation id generation algorithm like so:

```
@Begin(id="#{myConversationIdGenerator.nextId}")
public void editHotel() { ... }
```

Or it can be used to assign a meaningful conversation id:

```
@Begin(id="hotel#{hotel.id}")
public String editHotel() { ... }
```

```
@Begin(id="hotel#{hotelsDataModel.rowData.id}")
public String selectHotel() { ... }
```

```
@Begin(id="entry#{params['blogId']}")
public String viewBlogEntry() { ... }
```

```
@BeginTask(id="task#{taskInstance.id}")
public String approveDocument() { ... }
```

Clearly, these examples result in the same conversation id every time a particular hotel, blog or task is selected. So what happens if a conversation with the same conversation id already exists when the new conversation begins? Well, Seam detects the existing conversation and redirects to that conversation without running the `@Begin` method again. This feature helps control the number of workspaces that are created when using workspace management.

6.7. Natural conversation ids

When working with conversations that deal with persistent objects, it may be desirable to use the natural business key of the object instead of the standard, "surrogate" conversation id.

6.7.1. Configuration

Natural conversations are configured in `pages.xml`:

```
<conversation name="PlaceBid"
    parameter-name="auctionId"
    parameter-value="#{auction.auctionId}"/>
```

The first thing to note from the above configuration is that the conversation has a name, in this case `PlaceBid`. This name uniquely identifies this particular named conversation, and is used by the page definition to identify a named conversation to participate in.

The next attribute, `parameter-name` defines the request parameter that will contain the natural conversation id, in place of the default conversation id parameter (which is typically either `cid` or `conversationId`). In this example, the `parameter-name` is `auctionId`. This means that instead of a conversation parameter like `cid=123` appearing in the URL for your page, it will contain `auctionId=765432` instead.

The last attribute in the above configuration, `parameter-value`, defines an EL expression used to evaluate the value of the natural business key to use as the conversation id. In this example, the conversation id will be the primary key value of the auction instance currently in scope.

The next step is to configure which pages will participate in the named conversation. This is done by specifying the conversation attribute for a page definition:

```
<page view-id="/bid.xhtml" conversation="PlaceBid" login-required="true">
    <navigation from-action="#{bidAction.confirmBid}">
        <rule if-outcome="success">
            <redirect view-id="/auction.xhtml">
                <param name="id" value="#{bidAction.bid.auction.auctionId}"/>
            </redirect>
        </rule>
    </navigation>
</page>
```

6.7.2. Redirecting to a natural conversation

When starting, or redirecting to, a natural conversation there are a number of options for specifying the natural conversation name. Let's start by looking at the following page configuration:

```
<page view-id="/auction.xhtml">
    <param name="id" value="#{auctionDetail.selectedAuctionId}"/>

    <navigation from-action="#{bidAction.placeBid}">
        <redirect view-id="/bid.xhtml"/>
    </navigation>
</page>
```

```
</navigation>
</page>
```

From here, we can see that invoking the action `#{bidAction.placeBid}` from our auction view (by the way, all these examples are taken from the `seamBay` example in `Seam`), that we will be redirected to `/bid.xhtml`, which as we saw previously is configured with the natural conversation `PlaceBid`. The declaration for our action method looks like this:

```
@Begin(join = true)
public void placeBid()
```

When named conversations are specified in the `<page/>` element, redirection to the named conversation occurs as part of navigation rules, after the action method has already been invoked. This is a problem when redirecting to an existing conversation, as redirection needs to occur before the action method is invoked. To cater for this, it is necessary to specify the conversation name when the action is invoked. One way of doing this is by using the `s:conversationName` tag:

```
<h:commandButton id="placeBidWithAmount" styleClass="placeBid" action="#{bidAction.placeBid}">
  <s:conversationName value="PlaceBid"/>
</h:commandButton>
```

Another alternative is to specify the `conversationName` attribute when using either `s:link` or `s:button`:

```
<s:link value="Place Bid" action="#{bidAction.placeBid}" conversationName="PlaceBid"/>
```

6.8. Workspace management

Workspace management is the ability to "switch" conversations in a single window. `Seam` makes workspace management completely transparent at the level of the Java code. To enable workspace management, all you need to do is:

- Provide *description* text for each view id (when using JSF or `Seam` navigation rules) or page node (when using jPDL pageflows). This description text is displayed to the user by the workspace switchers.
- Include one or more of the standard workspace switcher JSP or facelets fragments in your pages. The standard fragments support workspace management via a drop down menu, a list of conversations, or breadcrumbs.

6.8.1. Workspace management and JSF navigation

When you use JSF or `Seam` navigation rules, `Seam` switches to a conversation by restoring the current `view-id` for that conversation. The descriptive text for the workspace is defined in a file called `pages.xml` that `Seam` expects to find in the `WEB-INF` directory, right next to `faces-config.xml`:

```
<pages>
  <page view-id="/main.xhtml">Search hotels: #{hotelBooking.searchString}</page>
  <page view-id="/hotel.xhtml">View hotel: #{hotel.name}</page>
  <page view-id="/book.xhtml">Book hotel: #{hotel.name}</page>
  <page view-id="/confirm.xhtml">Confirm: #{booking.description}</page>
</pages>
```

Note that if this file is missing, the `Seam` application will continue to work perfectly! The only missing functionality will be the ability to switch workspaces.

6.8.2. Workspace management and jPDL pageflow

When you use a jPDL pageflow definition, Seam switches to a conversation by restoring the current jBPM process state. This is a more flexible model since it allows the same `view-id` to have different descriptions depending upon the current `<page>` node. The description text is defined by the `<page>` node:

```
<pageflow-definition name="shopping">

  <start-state name="start">
    <transition to="browse"/>
  </start-state>

  <page name="browse" view-id="/browse.xhtml">
    <description>DVD Search: #{search.searchPattern}</description>
    <transition to="browse"/>
    <transition name="checkout" to="checkout"/>
  </page>

  <page name="checkout" view-id="/checkout.xhtml">
    <description>Purchase: ${cart.total}</description>
    <transition to="checkout"/>
    <transition name="complete" to="complete"/>
  </page>

  <page name="complete" view-id="/complete.xhtml">
    <end-conversation />
  </page>

</pageflow-definition>
```


6.8.3. The conversation switcher

Include the following fragment in your JSP or facelets page to get a drop-down menu that lets you switch to any current conversation, or to any other page of the application:

```
<h:selectOneMenu value="#{switcher.conversationIdOrOutcome}">
  <f:selectItem itemLabel="Find Issues" itemValue="findIssue"/>
  <f:selectItem itemLabel="Create Issue" itemValue="editIssue"/>
  <f:selectItems value="#{switcher.selectItems}"/>
</h:selectOneMenu>
<h:commandButton action="#{switcher.select}" value="Switch"/>
```

In this example, we have a menu that includes an item for each conversation, together with two additional items that let the user begin a new conversation.

Only conversations with a description (specified in `pages.xml`) will be included in the drop-down menu.



6.8.4. The conversation list

The conversation list is very similar to the conversation switcher, except that it is displayed as a table:

```
<h:dataTable value="#{conversationList}" var="entry"
  rendered="#{not empty conversationList}">
  <h:column>
    <f:facet name="header">Workspace</f:facet>
    <h:commandLink action="#{entry.select}" value="#{entry.description}"/>
    <h:outputText value="[current]" rendered="#{entry.current}"/>
  </h:column>
  <h:column>
    <f:facet name="header">Activity</f:facet>
    <h:outputText value="#{entry.startDatetime}">
      <f:convertDateTime type="time" pattern="hh:mm a"/>
    </h:outputText>
    <h:outputText value=" - "/>
    <h:outputText value="#{entry.lastDatetime}">
      <f:convertDateTime type="time" pattern="hh:mm a"/>
    </h:outputText>
  </h:column>
  <h:column>
    <f:facet name="header">Action</f:facet>
    <h:commandButton action="#{entry.select}" value="#{msg.Switch}"/>
    <h:commandButton action="#{entry.destroy}" value="#{msg.Destroy}"/>
  </h:column>
</h:dataTable>
```

We imagine that you will want to customize this for your own application.

Workspace	Workspace activity	Action
Comment on Issue [1] for Project [HHH]	01:18 PM - 01:18 PM	<button>Switch</button> <button>Destroy</button>
Issue [1] for Project [HHH]	01:18 PM - 01:18 PM	<button>Switch</button> <button>Destroy</button>
Project [HHH]	01:18 PM - 01:18 PM	<button>Switch</button> <button>Destroy</button>

Only conversations with a description will be included in the list.

Notice that the conversation list lets the user destroy workspaces.

6.8.5. Breadcrumbs

Breadcrumbs are useful in applications which use a nested conversation model. The breadcrumbs are a list of links to conversations in the current conversation stack:

```
<ui:repeat value="#{conversationStack}" var="entry">
  <h:outputText value=" | "/>
  <h:commandLink value="#{entry.description}" action="#{entry.select}"/>
</ui:repeat>
```

[Home](#) | [Find Issues](#) | [Create Issue](#) | [Project \[HHH\]](#) | [Issue \[1\] for Project \[HHH\]](#)

└─ Issue Attributes ─┘

6.9. Conversational components and JSF component bindings

Conversational components have one minor limitation: they cannot be used to hold bindings to JSF compon-

ents. (We generally prefer not to use this feature of JSF unless absolutely necessary, since it creates a hard dependency from application logic to the view.) On a postback request, component bindings are updated during the Restore View phase, before the Seam conversation context has been restored.

To work around this use an event scoped component to store the component bindings and inject it into the conversation scoped component that requires it.

```
@Name("grid")
@Scope(ScopeType.EVENT)
public class Grid
{
    private HtmlPanelGrid htmlPanelGrid;

    // getters and setters
    ...
}
```

```
@Name("gridEditor")
@Scope(ScopeType.CONVERSATION)
public class GridEditor
{
    @In(required=false)
    private Grid grid;

    ...
}
```

Alternatively, you can access the JSF component tree through the implicit `uiComponent` handle. The following example accesses `getRowIndex()` of the `UIData` component which backs the data table during iteration, it prints the current row number:

```
<h:dataTable id="lineItemTable" var="lineItem" value="#{orderHome.lineItems}">
  <h:column>
    Row: #{uiComponent['lineItemTable'].rowIndex}
  </h:column>
  ...
</h:dataTable>
```

JSF UI components are available with their client identifier in this map.

6.10. Concurrent calls to conversational components

A general discussion of concurrent calls to Seam components can be found in Section 3.1.10, “Concurrency model”. Here we will discuss the most common situation in which you will encounter concurrency — accessing conversational components from AJAX requests. We’re going to discuss the options that a Ajax client library should provide to control events originating at the client — and we’ll look at the options RichFaces gives you.

Conversational components don’t allow real concurrent access therefore Seam queues each request to process them serially. This allows each request to be executed in a deterministic fashion. However, a simple queue isn’t that great — firstly, if a method is, for some reason, taking a very long time to complete, running it over and over again whenever the client generates a request is bad idea (potential for Denial of Service attacks), and, secondly, AJAX is often to used to provide a quick status update to the user, so continuing to run the action after a long time isn’t useful.

Therefore Seam queues the action event for a period of time (the concurrent request timeout); if it can’t process

the event in time, it creates a temporary conversation and prints out a message to the user to let them know what's going on. It's therefore very important not to flood the server with AJAX events!

We can set a sensible default for the concurrent request timeout (in ms) in `components.xml`:

```
<core:manager concurrent-request-timeout="500" />
```

So far we've discussed "synchronous" AJAX requests - the client tells the server that an event has occur, and then rerenders part of the page based on the result. This approach is great when the AJAX request is lightweight (the methods called are simple e.g. calculating the sum of a column of numbers). But what if we need to do a complex computation?

For heavy computation we should use a truly asynchronous (poll based) approach — the client sends an AJAX request to the server, which causes action to be executed asynchronously on the server (so the the response to the client is immediate); the client then polls the server for updates. This is useful when you have a long-running action for which it is important that every action executes (you don't want some to be dropped as duplicates, or to timeout).

How should we design our conversational AJAX application?

Well first, you need to decide whether you want to use the simpler "synchronous" request or whether you want to add using a poll-style approach.

If you go for a "synchronous" approach, then you need to make an estimate of how long your AJAX request will take to complete - is it much shorter than the concurrent request timeout? If not, you probably want to alter the concurrent request timeout for this method (as discussed above). Next you probably want a queue on the client side to prevent flooding the server with requests. If the event occurs often (e.g. a keypress, onblur of input fields) and immediate update of the client is not a priority you should set a request delay on the client side. When working out your request delay, factor in that the event may also be queued on the server side.

Finally, the client library may provide an option to abort unfinished duplicate requests in favor of the most recent. You need to be careful with this option as it can lead to flooding of the server with requests if the server is not able to abort the unfinished request.

Using a poll-style design requires less fine-tuning. You just mark your action method `@Asynchronous` and decide on a polling interval:

```
int total;

// This method is called when an event occurs on the client
// It takes a really long time to execute
@Asynchronous
public void calculateTotal() {
    total = someReallyComplicatedCalculation();
}

// This method is called as the result of the poll
// It's very quick to execute
public int getTotal() {
    return total;
}
```

6.10.1. RichFaces Ajax

RichFaces Ajax is the AJAX library most commonly used with Seam, and provides all the controls discussed above:

- `eventsQueue` — provide a queue in which events are placed. All events are queued and requests are sent to the server serially. This is useful if the request can to the server can take some time to execute (e.g. heavy computation, retrieving information from a slow source) as the server isn't flooded.
- `ignoreDupResponses` — ignore the response produced by the request if a more recent 'similar' request is already in the queue. `ignoreDupResponses="true"` does *not cancel* the the processing of the request on the server side — just prevents unnecessary updates on the client side.

This option should be used with care with Seam's conversations as it allows multiple concurrent requests to be made.

- `requestDelay` — defines the time (in ms.) that the request will be remain on the queue. If the request has not been processed by after this time the request will be sent (regardless of whether a response has been received) or discarded (if there is a more recent similar event on the queue).

This option should be used with care with Seam's conversations as it allows multiple concurrent requests to be made. You need to be sure that the delay you set (in combination with the concurrent request timeout) is longer than the action will take to execute.

- `<a:poll reRender="total" interval="1000" />` — Polls the server, and rerenders an area as needed

Chapter 7. Pageflows and business processes

JBoss jBPM is a business process management engine for any Java SE or EE environment. jBPM lets you represent a business process or user interaction as a graph of nodes representing wait states, decisions, tasks, web pages, etc. The graph is defined using a simple, very readable, XML dialect called jPDL, and may be edited and visualised graphically using an eclipse plugin. jPDL is an extensible language, and is suitable for a range of problems, from defining web application page flow, to traditional workflow management, all the way up to orchestration of services in a SOA environment.

Seam applications use jBPM for two different problems:

- Defining the pageflow involved in complex user interactions. A jPDL process definition defines the page flow for a single conversation. A Seam conversation is considered to be a relatively short-running interaction with a single user.
- Defining the overarching business process. The business process may span multiple conversations with multiple users. Its state is persistent in the jBPM database, so it is considered long-running. Coordination of the activities of multiple users is a much more complex problem than scripting an interaction with a single user, so jBPM offers sophisticated facilities for task management and dealing with multiple concurrent paths of execution.

Don't get these two things confused ! They operate at very different levels or granularity. *Pageflow*, *conversation* and *task* all refer to a single interaction with a single user. A business process spans many tasks. Furthermore, the two applications of jBPM are totally orthogonal. You can use them together or independently or not at all.

You don't have to know jPDL to use Seam. If you're perfectly happy defining pageflow using JSF or Seam navigation rules, and if your application is more data-driven than process-driven, you probably don't need jBPM. But we're finding that thinking of user interaction in terms of a well-defined graphical representation is helping us build more robust applications.

7.1. Pageflow in Seam

There are two ways to define pageflow in Seam:

- Using JSF or Seam navigation rules - the *stateless navigation model*
- Using jPDL - the *stateful navigation model*

Very simple applications will only need the stateless navigation model. Very complex applications will use both models in different places. Each model has its strengths and weaknesses!

7.1.1. The two navigation models

The stateless model defines a mapping from a set of named, logical outcomes of an event directly to the resulting page of the view. The navigation rules are entirely oblivious to any state held by the application other than what page was the source of the event. This means that your action listener methods must sometimes make decisions about the page flow, since only they have access to the current state of the application.

Here is an example page flow definition using JSF navigation rules:

```

<navigation-rule>
  <from-view-id>/numberGuess.jsp</from-view-id>

  <navigation-case>
    <from-outcome>guess</from-outcome>
    <to-view-id>/numberGuess.jsp</to-view-id>
    <redirect/>
  </navigation-case>

  <navigation-case>
    <from-outcome>win</from-outcome>
    <to-view-id>/win.jsp</to-view-id>
    <redirect/>
  </navigation-case>

  <navigation-case>
    <from-outcome>lose</from-outcome>
    <to-view-id>/lose.jsp</to-view-id>
    <redirect/>
  </navigation-case>
</navigation-rule>

```

Here is the same example page flow definition using Seam navigation rules:

```

<page view-id="/numberGuess.jsp">
  <navigation>
    <rule if-outcome="guess">
      <redirect view-id="/numberGuess.jsp"/>
    </rule>
    <rule if-outcome="win">
      <redirect view-id="/win.jsp"/>
    </rule>
    <rule if-outcome="lose">
      <redirect view-id="/lose.jsp"/>
    </rule>
  </navigation>
</page>

```

If you find navigation rules overly verbose, you can return view ids directly from your action listener methods:

```

public String guess() {
    if (guess==randomNumber) return "/win.jsp";
    if (++guessCount==maxGuesses) return "/lose.jsp";
    return null;
}

```

Note that this results in a redirect. You can even specify parameters to be used in the redirect:

```

public String search() {
    return "/searchResults.jsp?searchPattern=#{searchAction.searchPattern}";
}

```

The stateful model defines a set of transitions between a set of named, logical application states. In this model, it is possible to express the flow of any user interaction entirely in the jPDL pageflow definition, and write action listener methods that are completely unaware of the flow of the interaction.

Here is an example page flow definition using jPDL:

```

<pageflow-definition name="numberGuess">
  <start-page name="displayGuess" view-id="/numberGuess.jsp">
    <redirect/>
  </start-page>
</pageflow-definition>

```

```

    <transition name="guess" to="evaluateGuess">
        <action expression="#{numberGuess.guess}" />
    </transition>
</start-page>

<decision name="evaluateGuess" expression="#{numberGuess.correctGuess}">
    <transition name="true" to="win"/>
    <transition name="false" to="evaluateRemainingGuesses"/>
</decision>


<decision name="evaluateRemainingGuesses" expression="#{numberGuess.lastGuess}">
    <transition name="true" to="lose"/>
    <transition name="false" to="displayGuess"/>
</decision>

<page name="win" view-id="/win.jsp">
    <redirect/>
    <end-conversation />
</page>

<page name="lose" view-id="/lose.jsp">
    <redirect/>
    <end-conversation />
</page>

</pageflow-definition>

```



There are two things we notice immediately here:

- The JSF/Seam navigation rules are *much* simpler. (However, this obscures the fact that the underlying Java code is more complex.)
- The jPDL makes the user interaction immediately understandable, without us needing to even look at the JSP or Java code.

In addition, the stateful model is more *constrained*. For each logical state (each step in the page flow), there are

a constrained set of possible transitions to other states. The stateless model is an *ad hoc* model which is suitable to relatively unconstrained, freeform navigation where the user decides where he/she wants to go next, not the application.

The stateful/stateless navigation distinction is quite similar to the traditional view of modal/modeless interaction. Now, Seam applications are not usually modal in the simple sense of the word - indeed, avoiding application modal behavior is one of the main reasons for having conversations! However, Seam applications can be, and often are, modal at the level of a particular conversation. It is well-known that modal behavior is something to avoid as much as possible; it is very difficult to predict the order in which your users are going to want to do things! However, there is no doubt that the stateful model has its place.

The biggest contrast between the two models is the back-button behavior.

7.1.2. Seam and the back button

When JSF or Seam navigation rules are used, Seam lets the user freely navigate via the back, forward and re-refresh buttons. It is the responsibility of the application to ensure that conversational state remains internally consistent when this occurs. Experience with the combination of web application frameworks like Struts or WebWork - that do not support a conversational model - and stateless component models like EJB stateless session beans or the Spring framework has taught many developers that this is close to impossible to do! However, our experience is that in the context of Seam, where there is a well-defined conversational model, backed by stateful session beans, it is actually quite straightforward. Usually it is as simple as combining the use of `no-conversation-view-id` with null checks at the beginning of action listener methods. We consider support for freeform navigation to be almost always desirable.

In this case, the `no-conversation-view-id` declaration goes in `pages.xml`. It tells Seam to redirect to a different page if a request originates from a page rendered during a conversation, and that conversation no longer exists:

```
<page view-id="/checkout.xhtml"
      no-conversation-view-id="/main.xhtml"/>
```

On the other hand, in the stateful model, backbuttoning is interpreted as an undefined transition back to a previous state. Since the stateful model enforces a defined set of transitions from the current state, back buttoning is by default disallowed in the stateful model! Seam transparently detects the use of the back button, and blocks any attempt to perform an action from a previous, "stale" page, and simply redirects the user to the "current" page (and displays a faces message). Whether you consider this a feature or a limitation of the stateful model depends upon your point of view: as an application developer, it is a feature; as a user, it might be frustrating! You can enable backbutton navigation from a particular page node by setting `back="enabled"`.

```
<page name="checkout"
      view-id="/checkout.xhtml"
      back="enabled">
  <redirect/>
  <transition to="checkout"/>
  <transition name="complete" to="complete"/>
</page>
```

This allows backbuttoning *from* the checkout state to *any previous state*!

Of course, we still need to define what happens if a request originates from a page rendered during a pageflow, and the conversation with the pageflow no longer exists. In this case, the `no-conversation-view-id` declaration goes into the pageflow definition:

```
<page name="checkout"
      view-id="/checkout.xhtml"
```

```

        back="enabled"
        no-conversation-view-id="/main.xhtml">
    <redirect/>
    <transition to="checkout"/>
    <transition name="complete" to="complete"/>
</page>

```

In practice, both navigation models have their place, and you'll quickly learn to recognize when to prefer one model over the other.

7.2. Using jPDL pageflows

7.2.1. Installing pageflows

We need to install the Seam jBPM-related components, and tell them where to find our pageflow definition. We can specify this Seam configuration in `components.xml`.

```

<bpm:jbp>
  <bpm:pageflow-definitions>
    <value>pageflow.jpdl.xml</value>
  </bpm:pageflow-definitions>
</bpm:jbp>

```

The first line installs jBPM, the second points to a jPDL-based pageflow definition.

7.2.2. Starting pageflows

We "start" a jPDL-based pageflow by specifying the name of the process definition using a `@Begin`, `@BeginTask` or `@StartTask` annotation:

```

@Begin(pageflow="numberguess")
public void begin() { ... }

```

Alternatively we can start a pageflow using `pages.xml`:

```

<page>
  <begin-conversation pageflow="numberguess"/>
</page>

```

If we are beginning the pageflow during the `RENDER_RESPONSE` phase—during a `@Factory` or `@Create` method, for example—we consider ourselves to be already at the page being rendered, and use a `<start-page>` node as the first node in the pageflow, as in the example above.

But if the pageflow is begun as the result of an action listener invocation, the outcome of the action listener determines which is the first page to be rendered. In this case, we use a `<start-state>` as the first node in the pageflow, and declare a transition for each possible outcome:

```

<pageflow-definition name="viewEditDocument">

  <start-state name="start">
    <transition name="documentFound" to="displayDocument"/>
    <transition name="documentNotFound" to="notFound"/>
  </start-state>

  <page name="displayDocument" view-id="/document.jsp">
    <transition name="edit" to="editDocument"/>
  </page>
</pageflow-definition>

```

```

    <transition name="done" to="main"/>
  </page>

  ...

  <page name="notFound" view-id="/404.jsp">
    <end-conversation/>
  </page>

</pageflow-definition>

```

7.2.3. Page nodes and transitions

Each `<page>` node represents a state where the system is waiting for user input:

```

<page name="displayGuess" view-id="/numberGuess.jsp">
  <redirect/>
  <transition name="guess" to="evaluateGuess">
    <action expression="#{numberGuess.guess}" />
  </transition>
</page>

```

The `view-id` is the JSF view id. The `<redirect/>` element has the same effect as `<redirect/>` in a JSF navigation rule: namely, a post-then-redirect behavior, to overcome problems with the browser's refresh button. (Note that Seam propagates conversation contexts over these browser redirects. So there is no need for a Ruby on Rails style "flash" construct in Seam!)

The transition name is the name of a JSF outcome triggered by clicking a command button or command link in `numberGuess.jsp`.

```

<h:commandButton type="submit" value="Guess" action="guess"/>

```

When the transition is triggered by clicking this button, jBPM will activate the transition action by calling the `guess()` method of the `numberGuess` component. Notice that the syntax used for specifying actions in the jPDL is just a familiar JSF EL expression, and that the transition action handler is just a method of a Seam component in the current Seam contexts. So we have exactly the same event model for jBPM events that we already have for JSF events! (The *One Kind of Stuff* principle.)

In the case of a null outcome (for example, a command button with no action defined), Seam will signal the transition with no name if one exists, or else simply redisplay the page if all transitions have names. So we could slightly simplify our example pageflow and this button:

```

<h:commandButton type="submit" value="Guess"/>

```

Would fire the following un-named transition:

```

<page name="displayGuess" view-id="/numberGuess.jsp">
  <redirect/>
  <transition to="evaluateGuess">
    <action expression="#{numberGuess.guess}" />
  </transition>
</page>

```

It is even possible to have the button call an action method, in which case the action outcome will determine the transition to be taken:

```

<h:commandButton type="submit" value="Guess" action="#{numberGuess.guess}"/>

```

```
<page name="displayGuess" view-id="/numberGuess.jsp">
  <transition name="correctGuess" to="win"/>
  <transition name="incorrectGuess" to="evaluateGuess"/>
</page>
```

However, this is considered an inferior style, since it moves responsibility for controlling the flow out of the pageflow definition and back into the other components. It is much better to centralize this concern in the pageflow itself.

7.2.4. Controlling the flow

Usually, we don't need the more powerful features of jPDL when defining pageflows. We do need the `<decision>` node, however:

```
<decision name="evaluateGuess" expression="#{numberGuess.correctGuess}">
  <transition name="true" to="win"/>
  <transition name="false" to="evaluateRemainingGuesses"/>
</decision>
```

A decision is made by evaluating a JSF EL expression in the Seam contexts.

7.2.5. Ending the flow

We end the conversation using `<end-conversation>` or `@End`. (In fact, for readability, use of *both* is encouraged.)

```
<page name="win" view-id="/win.jsp">
  <redirect/>
  <end-conversation/>
</page>
```

Optionally, we can end a task, specify a jBPM transition name. In this case, Seam will signal the end of the current task in the overarching business process.

```
<page name="win" view-id="/win.jsp">
  <redirect/>
  <end-task transition="success"/>
</page>
```

7.2.6. Pageflow composition

It is possible to compose pageflows and have one pageflow pause while another pageflow executes. The `<process-state>` node pauses the outer pageflow, and begins execution of a named pageflow:

```
<process-state name="cheat">
  <sub-process name="cheat"/>
  <transition to="displayGuess"/>
</process-state>
```

The inner flow begins executing at a `<start-state>` node. When it reaches an `<end-state>` node, execution of the inner flow ends, and execution of the outer flow resumes with the transition defined by the `<process-state>` element.

7.3. Business process management in Seam

A business process is a well-defined set of tasks that must be performed by users or software systems according to well-defined rules about *who* can perform a task, and *when* it should be performed. Seam's jBPM integration makes it easy to display lists of tasks to users and let them manage their tasks. Seam also lets the application store state associated with the business process in the `BUSINESS_PROCESS` context, and have that state made persistent via jBPM variables.

A simple business process definition looks much the same as a page flow definition (*One Kind of Stuff*), except that instead of <page> nodes, we have <task-node> nodes. In a long-running business process, the wait states are where the system is waiting for some user to log in and perform a task.


```
<process-definition name="todo">

  <start-state name="start">
    <transition to="todo"/>
  </start-state>

  <task-node name="todo">
    <task name="todo" description="#{todoList.description}">
      <assignment actor-id="#{actor.id}"/>
    </task>
    <transition to="done"/>
  </task-node>

  <end-state name="done"/>

</process-definition>
```



It is perfectly possible that we might have both jPDL business process definitions and jPDL pageflow definitions in the same project. If so, the relationship between the two is that a single <task> in a business process corresponds to a whole pageflow <pageflow-definition>

7.4. Using jPDL business process definitions

7.4.1. Installing process definitions

We need to install jBPM, and tell it where to find the business process definitions:

```
<bpm:jbpm>
  <bpm:process-definitions>
    <value>todo.jpdl.xml</value>
  </bpm:process-definitions>
</bpm:jbpm>
```

7.4.2. Initializing actor ids

We always need to know what user is currently logged in. jBPM "knows" users by their *actor id* and *group actor ids*. We specify the current actor ids using the built in Seam component named actor:

```
@In Actor actor;

public String login() {
    ...
    actor.setId( user.getUserName() );
    actor.getGroupActorIds().addAll( user.getGroupNames() );
    ...
}
```

7.4.3. Initiating a business process

To initiate a business process instance, we use the @CreateProcess annotation:

```
@CreateProcess(definition="todo")
public void createTodo() { ... }
```

Alternatively we can initiate a business process using pages.xml:

```
<page>
  <create-process definition="todo" />
</page>
```

7.4.4. Task assignment

When a process reaches a task node, task instances are created. These must be assigned to users or user groups. We can either hardcode our actor ids, or delegate to a Seam component:

```
<task name="todo" description="#{todoList.description}">
  <assignment actor-id="#{actor.id}"/>
</task>
```

In this case, we have simply assigned the task to the current user. We can also assign tasks to a pool:

```
<task name="todo" description="#{todoList.description}">
  <assignment pooled-actors="employees"/>
</task>
```

7.4.5. Task lists

Several built-in Seam components make it easy to display task lists. The `pooledTaskInstanceList` is a list of

pooled tasks that users may assign to themselves:

```
<h:dataTable value="#{pooledTaskInstanceList}" var="task">
  <h:column>
    <f:facet name="header">Description</f:facet>
    <h:outputText value="#{task.description}"/>
  </h:column>
  <h:column>
    <s:link action="#{pooledTask.assignToCurrentActor}" value="Assign" taskInstance="#{task}"/>
  </h:column>
</h:dataTable>
```

Note that instead of `<s:link>` we could have used a plain JSF `<h:commandLink>`:

```
<h:commandLink action="#{pooledTask.assignToCurrentActor}">
  <f:param name="taskId" value="#{task.id}"/>
</h:commandLink>
```

The `pooledTask` component is a built-in component that simply assigns the task to the current user.

The `taskInstanceListForType` component includes tasks of a particular type that are assigned to the current user:

```
<h:dataTable value="#{taskInstanceListForType['todo']}" var="task">
  <h:column>
    <f:facet name="header">Description</f:facet>
    <h:outputText value="#{task.description}"/>
  </h:column>
  <h:column>
    <s:link action="#{todoList.start}" value="Start Work" taskInstance="#{task}"/>
  </h:column>
</h:dataTable>
```

7.4.6. Performing a task

To begin work on a task, we use either `@StartTask` or `@BeginTask` on the listener method:

```
@StartTask
public String start() { ... }
```

Alternatively we can begin work on a task using `pages.xml`:

```
<page>
  <start-task />
</page>
```

These annotations begin a special kind of conversation that has significance in terms of the overarching business process. Work done by this conversation has access to state held in the business process context.

If we end the conversation using `@EndTask`, Seam will signal the completion of the task:

```
@EndTask(transition="completed")
public String completed() { ... }
```

Alternatively we can use `pages.xml`:

```
<page>
  <end-task transition="completed" />
</page>
```

You can also use EL to specify the transition in pages.xml.

At this point, jBPM takes over and continues executing the business process definition. (In more complex processes, several tasks might need to be completed before process execution can resume.)

Please refer to the jBPM documentation for a more thorough overview of the sophisticated features that jBPM provides for managing complex business processes.

Chapter 8. Seam and Object/Relational Mapping

Seam provides extensive support for the two most popular persistence architectures for Java: Hibernate3, and the Java Persistence API introduced with EJB 3.0. Seam's unique state-management architecture allows the most sophisticated ORM integration of any web application framework.

8.1. Introduction

Seam grew out of the frustration of the Hibernate team with the statelessness typical of the previous generation of Java application architectures. The state management architecture of Seam was originally designed to solve problems relating to persistence—in particular problems associated with *optimistic transaction processing*. Scalable online applications always use optimistic transactions. An atomic (database/JTA) level transaction should not span a user interaction unless the application is designed to support only a very small number of concurrent clients. But almost all interesting work involves first displaying data to a user, and then, slightly later, updating the same data. So Hibernate was designed to support the idea of a persistence context which spanned an optimistic transaction.

Unfortunately, the so-called "stateless" architectures that preceded Seam and EJB 3.0 had no construct for representing an optimistic transaction. So, instead, these architectures provided persistence contexts scoped to the atomic transaction. Of course, this resulted in many problems for users, and is the cause of the number one user complaint about Hibernate: the dreaded `LazyInitializationException`. What we need is a construct for representing an optimistic transaction in the application tier.

EJB 3.0 recognizes this problem, and introduces the idea of a stateful component (a stateful session bean) with an *extended persistence context* scoped to the lifetime of the component. This is a partial solution to the problem (and is a useful construct in and of itself) however there are two problems:

- The lifecycle of the stateful session bean must be managed manually via code in the web tier (it turns out that this is a subtle problem and much more difficult in practice than it sounds).
- Propagation of the persistence context between stateful components in the same optimistic transaction is possible, but tricky.

Seam solves the first problem by providing conversations, and stateful session bean components scoped to the conversation. (Most conversations actually represent optimistic transactions in the data layer.) This is sufficient for many simple applications (such as the Seam booking demo) where persistence context propagation is not needed. For more complex applications, with many loosely-interacting components in each conversation, propagation of the persistence context across components becomes an important issue. So Seam extends the persistence context management model of EJB 3.0, to provide conversation-scoped extended persistence contexts.

8.2. Seam managed transactions

EJB session beans feature declarative transaction management. The EJB container is able to start a transaction transparently when the bean is invoked, and end it when the invocation ends. If we write a session bean method that acts as a JSF action listener, we can do all the work associated with that action in one transaction, and be sure that it is committed or rolled back when we finish processing the action. This is a great feature, and all that is needed by some Seam applications.

However, there is a problem with this approach. A Seam application may not perform all data access for a request from a single method call to a session bean.

- The request might require processing by several loosely-coupled components, each of which is called independently from the web layer. It is common to see several or even many calls per request from the web layer to EJB components in Seam.
- Rendering of the view might require lazy fetching of associations.

The more transactions per request, the more likely we are to encounter atomicity and isolation problems when our application is processing many concurrent requests. Certainly, all write operations should occur in the same transaction!

Hibernate users developed the "*open session in view*" pattern to work around this problem. In the Hibernate community, "open session in view" was historically even more important because frameworks like Spring use transaction-scoped persistence contexts. So rendering the view would cause `LazyInitializationExceptions` when unfetched associations were accessed.

This pattern is usually implemented as a single transaction which spans the entire request. There are several problems with this implementation, the most serious being that we can never be sure that a transaction is successful until we commit it—but by the time the "open session in view" transaction is committed, the view is fully rendered, and the rendered response may already have been flushed to the client. How can we notify the user that their transaction was unsuccessful?

Seam solves both the transaction isolation problem and the association fetching problem, while working around the problems with "open session in view". The solution comes in two parts:

- use an extended persistence context that is scoped to the conversation, instead of to the transaction
- use two transactions per request; the first spans the beginning of the update model values phase until the end of the invoke application phase; the second spans the render response phase

In the next section, we'll tell you how to set up a conversation-scope persistence context. But first we need to tell you how to enable Seam transaction management. Note that you can use conversation-scoped persistence contexts without Seam transaction management, and there are good reasons to use Seam transaction management even when you're not using Seam-managed persistence contexts. However, the two facilities were designed to work together, and work best when used together.

Seam transaction management is useful even if you're using EJB 3.0 container-managed persistence contexts. But it is especially useful if you use Seam outside a Java EE 5 environment, or in any other case where you would use a Seam-managed persistence context.

8.2.1. Disabling Seam-managed transactions

Seam transaction management is enabled by default for all JSF requests. If you want to *disable* this feature, you can do it in `components.xml`:

```
<core:init transaction-management-enabled="false"/>
<transaction:no-transaction />
```

8.2.2. Configuring a Seam transaction manager

Seam provides a transaction management abstraction for beginning, committing, rolling back, and synchronizing with a transaction. By default Seam uses a JTA transaction component that integrates with Container Managed and programmatic EJB transactions.

Out of the box Seam provides transaction components for the following additional transaction APIs:

- JPA RESOURCE_LOCAL transactions with the `javax.persistence.EntityTransaction` interface.
- Hibernate managed transactions with the `org.hibernate.Transaction` interface.
- Spring managed transactions with the `org.springframework.transaction.PlatformTransactionManager` interface.
- Explicitly disable Seam managed transactions

Configure JPA RESOURCE_LOCAL transaction management by adding the following to your `components.xml` where `#{em}` is the name of the `persistence:managed-persistence-context` component. If your managed persistence context is named `entityManager`, you can opt to leave out the `entity-manager` attribute. (see Seam-managed persistence contexts)

```
<transaction:entity-transaction entity-manager="#{em}"/>
```

To configure Hibernate managed transactions declare the following in your `components.xml` where `#{hibernateSession}` is the name of the project's `persistence:managed-hibernate-session` component. If your managed hibernate session is named `session`, you can opt to leave out the `session` attribute. (see Seam-managed persistence contexts)

```
<transaction:hibernate-transaction session="#{hibernateSession}"/>
```

To explicitly disable Seam managed transactions declare the following in your `components.xml`:

```
<transaction:no-transaction />
```

For configuring Spring managed transactions see using Spring `PlatformTransactionManagement` .

8.2.3. Transaction synchronization

Transaction synchronization provides callbacks for transaction related events such as `beforeCompletion()` and `afterCompletion()`. By default, Seam uses it's own transaction synchronization component which requires explicit use of the Seam transaction component when committing a transaction to ensure synchronization callbacks are correctly executed. If in a Java EE 5 environment the `<transaction:ejb-transaction/>` component should be declared in `components.xml` to ensure that Seam synchronization callbacks are correctly called if the container commits a transaction outside of Seam's knowledge.

8.3. Seam-managed persistence contexts

If you're using Seam outside of a Java EE 5 environment, you can't rely upon the container to manage the persistence context lifecycle for you. Even if you are in an EE 5 environment, you might have a complex application with many loosely coupled components that collaborate together in the scope of a single conversation, and in this case you might find that propagation of the persistence context between component is tricky and error-prone.

In either case, you'll need to use a *managed persistence context* (for JPA) or a *managed session* (for Hibernate) in your components. A Seam-managed persistence context is just a built-in Seam component that manages an instance of `EntityManager` or `Session` in the conversation context. You can inject it with `@In`.

Seam-managed persistence contexts are extremely efficient in a clustered environment. Seam is able to perform an optimization that EJB 3.0 specification does not allow containers to use for container-managed extended persistence contexts. Seam supports transparent failover of extended persistence contexts, without the need to replicate any persistence context state between nodes. (We hope to fix this oversight in the next revision of the EJB spec.)

8.3.1. Using a Seam-managed persistence context with JPA

Configuring a managed persistence context is easy. In `components.xml`, we can write:

```
<persistence:managed-persistence-context name="bookingDatabase"
    auto-create="true"
    persistence-unit-jndi-name="java:/EntityManagerFactories/bookingData"/>
```

This configuration creates a conversation-scoped Seam component named `bookingDatabase` that manages the lifecycle of `EntityManager` instances for the persistence unit (`EntityManagerFactory` instance) with JNDI name `java:/EntityManagerFactories/bookingData`.

Of course, you need to make sure that you have bound the `EntityManagerFactory` into JNDI. In JBoss, you can do this by adding the following property setting to `persistence.xml`.

```
<property name="jboss.entity.manager.factory.jndi.name"
    value="java:/EntityManagerFactories/bookingData"/>
```

Now we can have our `EntityManager` injected using:

```
@In EntityManager bookingDatabase;
```

If you are using EJB3 and mark your class or method `@TransactionAttribute(REQUIRES_NEW)` then the transaction and persistence context shouldn't be propagated to method calls on this object. However as the Seam-managed persistence context is propagated to any component within the conversation, it will be propagated to methods marked `REQUIRES_NEW`. Therefore, if you mark a method `REQUIRES_NEW` then you should access the entity manager using `@PersistenceContext`.

8.3.2. Using a Seam-managed Hibernate session

Seam-managed Hibernate sessions are similar. In `components.xml`:

```
<persistence:hibernate-session-factory name="hibernateSessionFactory"/>
<persistence:managed-hibernate-session name="bookingDatabase"
    auto-create="true"
    session-factory-jndi-name="java:/bookingSessionFactory"/>
```

Where `java:/bookingSessionFactory` is the name of the session factory specified in `hibernate.cfg.xml`.

```
<session-factory name="java:/bookingSessionFactory">
  <property name="transaction.flush_before_completion">true</property>
  <property name="connection.release_mode">after_statement</property>
  <property name="transaction.manager_lookup_class">org.hibernate.transaction.JBossTransactionManager</property>
  <property name="transaction.factory_class">org.hibernate.transaction.JTATransactionFactory</property>
  <property name="connection.datasource">java:/bookingDatasource</property>
```

```
...
</session-factory>
```

Note that Seam does not flush the session, so you should always enable `hibernate.transaction.flush_before_completion` to ensure that the session is automatically flushed before the JTA transaction commits.

We can now have a managed Hibernate Session injected into our JavaBean components using the following code:

```
@In Session bookingDatabase;
```

8.3.3. Seam-managed persistence contexts and atomic conversations

Persistence contexts scoped to the conversation allows you to program optimistic transactions that span multiple requests to the server without the need to use the `merge()` operation, without the need to re-load data at the beginning of each request, and without the need to wrestle with the `LazyInitializationException` or `NonUniqueObjectException`.

As with any optimistic transaction management, transaction isolation and consistency can be achieved via use of optimistic locking. Fortunately, both Hibernate and EJB 3.0 make it very easy to use optimistic locking, by providing the `@Version` annotation.

By default, the persistence context is flushed (synchronized with the database) at the end of each transaction. This is sometimes the desired behavior. But very often, we would prefer that all changes are held in memory and only written to the database when the conversation ends successfully. This allows for truly atomic conversations. As the result of a truly stupid and shortsighted decision by certain non-JBoss, non-Sun and non-Sybase members of the EJB 3.0 expert group, there is currently no simple, usable and portable way to implement atomic conversations using EJB 3.0 persistence. However, Hibernate provides this feature as a vendor extension to the `FlushModeTypes` defined by the specification, and it is our expectation that other vendors will soon provide a similar extension.

Seam lets you specify `FlushModeType.MANUAL` when beginning a conversation. Currently, this works only when Hibernate is the underlying persistence provider, but we plan to support other equivalent vendor extensions.

```
@In EntityManager em; //a Seam-managed persistence context

@Begin(flushMode=MANUAL)
public void beginClaimWizard() {
    claim = em.find(Claim.class, claimId);
}
```

Now, the `claim` object remains managed by the persistence context for the rest of the conversation. We can make changes to the claim:

```
public void addPartyToClaim() {
    Party party = ....;
    claim.addParty(party);
}
```

But these changes will not be flushed to the database until we explicitly force the flush to occur:

```
@End
public void commitClaim() {
    em.flush();
}
```

Of course, you could set the `flushMode` to `MANUAL` from `pages.xml`, for example in a navigation rule:

```
<begin-conversation flush-mode="MANUAL" />
```

8.4. Using the JPA "delegate"

The `EntityManager` interface lets you access a vendor-specific API via the `getDelegate()` method. Naturally, the most interesting vendor is Hibernate, and the most powerful delegate interface is `org.hibernate.Session`. You'd be nuts to use anything else. Trust me, I'm not biased at all.

But regardless of whether you're using Hibernate (genius!) or something else (masochist, or just not very bright), you'll almost certainly want to use the delegate in your Seam components from time to time. One approach would be the following:

```
@In EntityManager entityManager;

@Create
public void init() {
    ( (Session) entityManager.getDelegate() ).enableFilter("currentVersions");
}
```

But typecasts are unquestionably the ugliest syntax in the Java language, so most people avoid them whenever possible. Here's a different way to get at the delegate. First, add the following line to `components.xml`:

```
<factory name="session"
        scope="STATELESS"
        auto-create="true"
        value="#{entityManager.delegate}"/>
```

Now we can inject the session directly:

```
@In Session session;

@Create
public void init() {
    session.enableFilter("currentVersions");
}
```

8.5. Using EL in EJB-QL/HQL

Seam proxies the `EntityManager` or `Session` object whenever you use a Seam-managed persistence context or inject a container managed persistence context using `@PersistenceContext`. This lets you use EL expressions in your query strings, safely and efficiently. For example, this:

```
User user = em.createQuery("from User where username=#{user.username}")
               .getSingleResult();
```

is equivalent to:

```
User user = em.createQuery("from User where username=:username")
               .setParameter("username", user.getUsername())
               .getSingleResult();
```

Of course, you should never, ever write it like this:

```
User user = em.createQuery("from User where username=" + user.getUsername()) //BAD!
    .getSingleResult();
```

(It is inefficient and vulnerable to SQL injection attacks.)

8.6. Using Hibernate filters

The coolest, and most unique, feature of Hibernate is *filters*. Filters let you provide a restricted view of the data in the database. You can find out more about filters in the Hibernate documentation. But we thought we'd mention an easy way to incorporate filters into a Seam application, one that works especially well with the Seam Application Framework.

Seam-managed persistence contexts may have a list of filters defined, which will be enabled whenever an EntityManager or Hibernate Session is first created. (Of course, they may only be used when Hibernate is the underlying persistence provider.)

```
<persistence:filter name="regionFilter">
  <persistence:name>region</persistence:name>
  <persistence:parameters>
    <key>regionCode</key>
    <value>#{region.code}</value>
  </persistence:parameters>
</persistence:filter>

<persistence:filter name="currentFilter">
  <persistence:name>current</persistence:name>
  <persistence:parameters>
    <key>date</key>
    <value>#{currentDate}</value>
  </persistence:parameters>
</persistence:filter>

<persistence:managed-persistence-context name="personDatabase"
  persistence-unit-jndi-name="java:/EntityManagerFactories/personDatabase">
  <core:filters>
    <value>#{regionFilter}</value>
    <value>#{currentFilter}</value>
  </core:filters>
</persistence:managed-persistence-context>
```

Chapter 9. JSF form validation in Seam

In plain JSF, validation is defined in the view:

```
<h:form>
  <h:messages/>

  <div>
    Country:
    <h:inputText value="#{location.country}" required="true">
      <my:validateCountry/>
    </h:inputText>
  </div>

  <div>
    Zip code:
    <h:inputText value="#{location.zip}" required="true">
      <my:validateZip/>
    </h:inputText>
  </div>

  <h:commandButton/>
</h:form>
```

In practice, this approach usually violates DRY, since most "validation" actually enforces constraints that are part of the data model, and exist all the way down to the database schema definition. Seam provides support for model-based constraints defined using Hibernate Validator.

Let's start by defining our constraints, on our Location class:

```
public class Location {
    private String country;
    private String zip;

    @NotNull
    @Length(max=30)
    public String getCountry() { return country; }
    public void setCountry(String c) { country = c; }

    @NotNull
    @Length(max=6)
    @Pattern("^\\d*$")
    public String getZip() { return zip; }
    public void setZip(String z) { zip = z; }
}
```

Well, that's a decent first cut, but in practice it might be more elegant to use custom constraints instead of the ones built into Hibernate Validator:

```
public class Location {
    private String country;
    private String zip;

    @NotNull
    @Country
    public String getCountry() { return country; }
    public void setCountry(String c) { country = c; }

    @NotNull
    @ZipCode
    public String getZip() { return zip; }
    public void setZip(String z) { zip = z; }
}
```

Whichever route we take, we no longer need to specify the type of validation to be used in the JSF page. Instead, we can use `<s:validate>` to validate against the constraint defined on the model object.

```
<h:form>
  <h:messages/>

  <div>
    Country:
    <h:inputText value="#{location.country}" required="true">
      <s:validate/>
    </h:inputText>
  </div>

  <div>
    Zip code:
    <h:inputText value="#{location.zip}" required="true">
      <s:validate/>
    </h:inputText>
  </div>

  <h:commandButton/>
</h:form>
```

Note: specifying `@NotNull` on the model does *not* eliminate the requirement for `required="true"` to appear on the control! This is due to a limitation of the JSF validation architecture.

This approach *defines* constraints on the model, and *presents* constraint violations in the view—a significantly better design.

However, it is not much less verbose than what we started with, so let's try `<s:validateAll>`:

```
<h:form>
  <h:messages/>

  <s:validateAll>
    <div>
      Country:
      <h:inputText value="#{location.country}" required="true"/>
    </div>

    <div>
      Zip code:
      <h:inputText value="#{location.zip}" required="true"/>
    </div>

    <h:commandButton/>
  </s:validateAll>
</h:form>
```

This tag simply adds an `<s:validate>` to every input in the form. For a large form, it can save a lot of typing!

Now we need to do something about displaying feedback to the user when validation fails. Currently we are displaying all messages at the top of the form. What we would really like to do is display the message next to the field with the error (this is possible in plain JSF), highlight the field and label (this is not possible) and, for good measure, display some image next the the field (also not possible). We also want to display a little colored asterisk next to the label for each required form field.

That's quite a lot of functionality we need for each field of our form. We wouldn't want to have to specify hig-

lighting and the layout of the image, message and input field for every field on the form. So, instead, we'll specify the common layout in a facelets template:

```
<ui:composition xmlns="http://www.w3.org/1999/xhtml"
    xmlns:ui="http://java.sun.com/jsf/facelets"
    xmlns:h="http://java.sun.com/jsf/html"
    xmlns:f="http://java.sun.com/jsf/core"
    xmlns:s="http://jboss.com/products/seam/taglib">

    <div>

        <s:label styleClass="#{invalid?'error':''}">
            <ui:insert name="label"/>
            <s:span styleClass="required" rendered="#{required}">*</s:span>
        </s:label>

        <span class="#{invalid?'error':''}">
            <h:graphicImage value="/img/error.gif" rendered="#{invalid}" />
            <s:validateAll>
                <ui:insert />
            </s:validateAll>
        </span>

        <s:message styleClass="error" />

    </div>

</ui:composition>
```

We can include this template for each of our form fields using `<s:decorate>`.

```
<h:form>

    <h:messages globalOnly="true" />

    <s:decorate template="edit.xhtml">
        <ui:define name="label">Country:</ui:define>
        <h:inputText value="#{location.country}" required="true" />
    </s:decorate>

    <s:decorate template="edit.xhtml">
        <ui:define name="label">Zip code:</ui:define>
        <h:inputText value="#{location.zip}" required="true" />
    </s:decorate>

    <h:commandButton />

</h:form>
```

Finally, we can use RichFaces Ajax to display validation messages as the user is navigating around the form:

```
<h:form>

    <h:messages globalOnly="true" />

    <s:decorate id="countryDecoration" template="edit.xhtml">
        <ui:define name="label">Country:</ui:define>
        <h:inputText value="#{location.country}" required="true">
            <a:support event="onblur" reRender="countryDecoration" bypassUpdates="true" />
        </h:inputText>
    </s:decorate>

    <s:decorate id="zipDecoration" template="edit.xhtml">
        <ui:define name="label">Zip code:</ui:define>
        <h:inputText value="#{location.zip}" required="true">
            <a:support event="onblur" reRender="zipDecoration" bypassUpdates="true" />
        </h:inputText>
    </s:decorate>

</h:form>
```

```

    <h:commandButton/>
</h:form>

```

It's better style to define explicit ids for important controls on the page, especially if you want to do automated testing for the UI, using some toolkit like Selenium. If you don't provide explicit ids, JSF will generate them, but the generated values will change if you change anything on the page.

```

<h:form id="form">

    <h:messages globalOnly="true"/>

    <s:decorate id="countryDecoration" template="edit.xhtml">
        <ui:define name="label">Country:</ui:define>
        <h:inputText id="country" value="#{location.country}" required="true">
            <a:support event="onblur" reRender="countryDecoration" bypassUpdates="true"/>
        </h:inputText>
    </s:decorate>

    <s:decorate id="zipDecoration" template="edit.xhtml">
        <ui:define name="label">Zip code:</ui:define>
        <h:inputText id="zip" value="#{location.zip}" required="true">
            <a:support event="onblur" reRender="zipDecoration" bypassUpdates="true"/>
        </h:inputText>
    </s:decorate>

    <h:commandButton/>

</h:form>

```

And what if you want to specify a different message to be displayed when validation fails? You can use the Seam message bundle (and all its goodies like el expressions inside the message, and per-view message bundles) with the Hibernate Validator:

```

public class Location {
    private String name;
    private String zip;

    // Getters and setters for name

    @NotNull
    @Length(max=6)
    @ZipCode(message="#{messages['location.zipCode.invalid']}")
    public String getZip() { return zip; }
    public void setZip(String z) { zip = z; }
}

```

```

location.zipCode.invalid = The zip code is not valid for #{location.name}

```

Chapter 10. Groovy integration

One aspect of JBoss Seam is its RAD (Rapid Application Development) capability. While not synonymous with RAD, one interesting tool in this space is dynamic languages. Until recently, choosing a dynamic language was required choosing a completely different development platform (a development platform with a set of APIs and a runtime so great that you would no longer want to use you old legacy Java [sic] APIs anymore, which would be lucky because you would be forced to use those proprietary APIs anyway). Dynamic languages built on top of the Java Virtual Machine, and Groovy [<http://groovy.codehaus.org>] in particular broke this approach in silos.

JBoss Seam now unites the dynamic language world with the Java EE world by seamlessly integrating both static and dynamic languages. JBoss Seam lets the application developer use the best tool for the task, without context switching. Writing dynamic Seam components is exactly like writing regular Seam components. You use the same annotations, the same APIs, the same everything.

10.1. Groovy introduction

Groovy is an agile dynamic language based on the Java language but with additional features inspired by Python, Ruby and Smalltalk. The strengths of Groovy are twofold:

- Java syntax is supported in Groovy: Java code is Groovy code, making the learning curve very smooth
- Groovy objects are Java objects, and Groovy classes are Java classes: Groovy integrates smoothly with existing Java libraries and frameworks.

TODO: write a quick overview of the Groovy syntax add-on

10.2. Writing Seam applications in Groovy

There is not much to say about it. Since a Groovy object is a Java object, you can virtually write any Seam component, or any class for what it worth, in Groovy and deploy it. You can also mix Groovy classes and Java classes in the same application.

10.2.1. Writing Groovy components

As you should have noticed by now, Seam uses annotations heavily. Be sure to use Groovy 1.1 Beta1 or above for annotation support. Here are some example of groovy code used in a Seam application.

10.2.1.1. Entity

```
@Entity
@Name("hotel")
class Hotel implements Serializable
{
    @Id @GeneratedValue
    Long id

    @Length(max=50) @NotNull
    String name

    @Length(max=100) @NotNull
    String address
}
```

```

    @Length(max=40) @NotNull
    String city

    @Length(min=2, max=10) @NotNull
    String state

    @Length(min=4, max=6) @NotNull
    String zip

    @Length(min=2, max=40) @NotNull
    String country

    @Column(precision=6, scale=2)
    BigDecimal price

    @Override
    String toString()
    {
        return "Hotel(${name},${address},${city},${zip})"
    }
}

```

Groovy natively support the notion of properties (getter/setter), so there is no need to explicitly write verbose getters and setters: in the previous example, the hotel class can be accessed from Java as `hotel.getCity()`, the getters and setters being generated by the Groovy compiler. This type of syntactic sugar makes the entity code very concise.

Some temporary gotchas: Groovy 1.1 Beta1 does not (yet) support generics. One negative effect of this is that entity relationships have no built-in type information. It is necessary to use the `targetEntity` attribute of the appropriate `OneToMany` annotation instead of simply using a generic type definition like `Collection<Entity>`. For the same reason, you won't be able to benefit from the very useful Chapter 11, *The Seam Application Framework* yet. The good news is that support for generics is targeted for Groovy 1.1 (Groovy 1.1 Beta2 at the time of writing).

10.2.1.2. Seam component

Writing Seam components in Groovy is in no way different than in Java: annotations are used to mark the class as a Seam component.

```

@Scope(ScopeType.SESSION)
@Name("bookingList")
class BookingListAction implements Serializable
{
    @In EntityManager em
    @In User user
    @DataModel List<Booking> bookings
    @DataModelSelection Booking booking
    @Logger Log log

    @Factory public void getBookings()
    {
        bookings = em.createQuery('''
            select b from Booking b
            where b.user.username = :username
            order by b.checkinDate''')
            .setParameter("username", user.username)
            .getResultList()
    }

    public void cancel()
    {
        log.info("Cancel booking: #{bookingList.booking.id} for #{user.username}")
        Booking cancelled = em.find(Booking.class, booking.id)
        if (cancelled != null) em.remove( cancelled )
    }
}

```

```
        getBookings()  
        FacesMessages.instance().add("Booking cancelled for confirmation number #{bookingList.booking  
    }  
}
```

10.2.2. seam-gen

Seam gen has a transparent integration with Groovy. You can write Groovy code in seam-gen backed projects without any additional infrastructure requirement. When writing a Groovy entity, simply place your `.groovy` files in `src/model`. Unsurprisingly, when writing an action, simply place your `.groovy` files in `src/action`.

10.3. Deployment

Deploying Groovy classes is very much like deploying Java classes (surprisingly, no need to write nor comply with a 3-letter composite specification to support a multi-language component framework).

Beyond standard deployments, JBoss Seam has the ability, at development time, to redeploy JavaBeans Seam component classes without having to restart the application, saving a lot of time in the development / test cycle. The same support is provided for GroovyBeans Seam components when the `.groovy` files are deployed.

10.3.1. Deploying Groovy code

A Groovy class *is* a Java class, with a bytecode representation just like a Java class. To deploy, a Groovy entity, a Groovy Session bean or a Groovy Seam component, a compilation step is necessary. A common approach is to use the `groovyc` ant task. Once compiled, a Groovy class is in no way different than a Java class and the application server will treat them equally. Note that this allow a seamless mix of Groovy and Java code.

10.3.2. Native `.groovy` file deployment at development time

JBoss Seam natively supports the deployment of `.groovy` files (ie without compilation) in incremental hotdeployment mode (development only). This enables a very fast edit/test cycle. To set up `.groovy` deployments, follow the configuration at Section 2.8, “Seam and incremental hot deployment” and deploy your Groovy code (`.groovy` files) into the `WEB-INF/dev` directory. The GroovyBean components will be picked up incrementally with no need to restart the application (and obviously not the application server either).

Be aware that the native `.groovy` file deployment suffers the same limitations as the regular Seam hotdeployment:

- The components must be JavaBeans or GroovyBeans. They cannot be EJB3 bean
- Entities cannot be hotdeployed
- The hot-deployable components will not be visible to any classes deployed outside of `WEB-INF/dev`
- Seam debug mode must be enabled

10.3.3. seam-gen

Seam-gen transparently supports Groovy files deployment and compilation. This includes the native `.groovy` file deployment in development mode (compilation-less). If you create a seam-gen project of type WAR, Java and Groovy classes in `src/action` will automatically be candidate for the incremental hot deployment. If you are in production mode, the Groovy files will simply be compiled before deployment.

You will find a live example of the Booking demo written completely in Groovy and supporting incremental hot deployment in `examples/groovybooking`.

Chapter 11. The Seam Application Framework

Seam makes it really easy to create applications by writing plain Java classes with annotations, which don't need to extend any special interfaces or superclasses. But we can simplify some common programming tasks even further, by providing a set of pre-built components which can be re-used either by configuration in `components.xml` (for very simple cases) or extension.

The *Seam Application Framework* can reduce the amount of code you need to write when doing basic database access in a web application, using either Hibernate or JPA.

We should emphasize that the framework is extremely simple, just a handful of simple classes that are easy to understand and extend. The "magic" is in Seam itself—the same magic you use when creating any Seam application even without using this framework.

11.1. Introduction

The components provided by the Seam application framework may be used in one of two different approaches. The first way is to install and configure an instance of the component in `components.xml`, just like we have done with other kinds of built-in Seam components. For example, the following fragment from `components.xml` installs a component which can perform basic CRUD operations for a `Person` entity:

```
<framework:entity-home name="personHome"
                        entity-class="eg.Person"
                        entity-manager="#{personDatabase}">
  <framework:id>#{param.personId}</framework:id>
</framework:entity-home>
```

If that looks a bit too much like "programming in XML" for your taste, you can use extension instead:

```
@Stateful
@Name("personHome")
public class PersonHome extends EntityHome<Person> implements LocalPersonHome {
    @RequestParameter String personId;
    @In EntityManager personDatabase;

    public Object getId() { return personId; }
    public EntityManager getEntityManager() { return personDatabase; }
}
```

The second approach has one huge advantage: you can easily add extra functionality, and override the built-in functionality (the framework classes were carefully designed for extension and customization).

A second advantage is that your classes may be EJB stateful session beans, if you like. (They do not have to be, they can be plain JavaBean components if you prefer.) If you are using JBoss AS, you'll need 4.2.2.GA or later.

At this time, the Seam Application Framework provides just four built-in components: `EntityHome` and `HibernateEntityHome` for CRUD, along with `EntityQuery` and `HibernateEntityQuery` for queries.

The Home and Query components are written so that they can function with a scope of session, event or conversation. Which scope you use depends upon the state model you wish to use in your application.

The Seam Application Framework only works with Seam-managed persistence contexts. By default, the components will look for a persistence context named `entityManager`.

11.2. Home objects

A Home object provides persistence operations for a particular entity class. Suppose we have our trusty Person class:

```
@Entity
public class Person {
    @Id private Long id;
    private String firstName;
    private String lastName;
    private Country nationality;

    //getters and setters...
}
```

We can define a personHome component either via configuration:

```
<framework:entity-home name="personHome" entity-class="eg.Person" />
```

Or via extension:

```
@Name("personHome")
public class PersonHome extends EntityHome<Person> {}
```

A Home object provides the following operations: `persist()`, `remove()`, `update()` and `getInstance()`. Before you can call the `remove()`, or `update()` operations, you must first set the identifier of the object you are interested in, using the `setId()` method.

We can use a Home directly from a JSF page, for example:

```
<h1>Create Person</h1>
<h:form>
    <div>First name: <h:inputText value="#{personHome.instance.firstName}"/></div>
    <div>Last name: <h:inputText value="#{personHome.instance.lastName}"/></div>
    <div>
        <h:commandButton value="Create Person" action="#{personHome.persist}"/>
    </div>
</h:form>
```

Usually, it is much nicer to be able to refer to the Person merely as `person`, so let's make that possible by adding a line to `components.xml`:

```
<factory name="person"
        value="#{personHome.instance}"/>

<framework:entity-home name="personHome"
        entity-class="eg.Person" />
```

(If we are using configuration.) Or by adding a `@Factory` method to `PersonHome`:

```
@Name("personHome")
public class PersonHome extends EntityHome<Person> {

    @Factory("person")
    public Person initPerson() { return getInstance(); }

}
```

(If we are using extension.) This change simplifies our JSF page to the following:

```

<h1>Create Person</h1>
<h:form>
    <div>First name: <h:inputText value="#{person.firstName}"/></div>
    <div>Last name: <h:inputText value="#{person.lastName}"/></div>
    <div>
        <h:commandButton value="Create Person" action="#{personHome.persist}"/>
    </div>
</h:form>

```

Well, that lets us create new Person entries. Yes, that is all the code that is required! Now, if we want to be able to display, update and delete pre-existing Person entries in the database, we need to be able to pass the entry identifier to the PersonHome. Page parameters are a great way to do that:

```

<pages>
    <page view-id="/editPerson.jsp">
        <param name="personId" value="#{personHome.id}"/>
    </page>
</pages>

```

Now we can add the extra operations to our JSF page:

```

<h1>
    <h:outputText rendered="#{!personHome.managed}" value="Create Person"/>
    <h:outputText rendered="#{personHome.managed}" value="Edit Person"/>
</h1>
<h:form>
    <div>First name: <h:inputText value="#{person.firstName}"/></div>
    <div>Last name: <h:inputText value="#{person.lastName}"/></div>
    <div>
        <h:commandButton value="Create Person" action="#{personHome.persist}" rendered="#{!personHome.managed}" />
        <h:commandButton value="Update Person" action="#{personHome.update}" rendered="#{personHome.managed}" />
        <h:commandButton value="Delete Person" action="#{personHome.remove}" rendered="#{personHome.managed}" />
    </div>
</h:form>

```

When we link to the page with no request parameters, the page will be displayed as a "Create Person" page. When we provide a value for the personId request parameter, it will be an "Edit Person" page.

Suppose we need to create Person entries with their nationality initialized. We can do that easily, via configuration:

```

<factory name="person"
    value="#{personHome.instance}"/>

<framework:entity-home name="personHome"
    entity-class="eg.Person"
    new-instance="#{newPerson}"/>

<component name="newPerson"
    class="eg.Person">
    <property name="nationality">#{country}</property>
</component>

```

Or by extension:

```

@Name("personHome")
public class PersonHome extends EntityHome<Person> {

    @In Country country;

    @Factory("person")
    public Person initPerson() { return getInstance(); }

    protected Person createInstance() {

```

```

        return new Person(country);
    }
}

```

Of course, the Country could be an object managed by another Home object, for example, CountryHome.

To add more sophisticated operations (association management, etc), we can just add methods to PersonHome.

```

@Name("personHome")
public class PersonHome extends EntityHome<Person> {

    @In Country country;

    @Factory("person")
    public Person initPerson() { return getInstance(); }

    protected Person createInstance() {
        return new Person(country);
    }

    public void migrate()
    {
        getInstance().setCountry(country);
        update();
    }
}

```

The Home object raises an `org.jboss.seam.afterTransactionSuccess` event when a transaction succeeds (a call to `persist()`, `update()` or `remove()` succeeds). By observing this event we can refresh our queries when the underlying entities are changed. If we only want to refresh certain queries when a particular entity is persisted, updated or removed we can observe the `org.jboss.seam.afterTransactionSuccess.<name>` event (where `<name>` is the name of the entity).

The Home object automatically displays faces messages when an operation is successful. To customize these messages we can, again, use configuration:

```

<factory name="person"
        value="#{personHome.instance}"/>

<framework:entity-home name="personHome"
                      entity-class="eg.Person"
                      new-instance="#{newPerson}">
    <framework:created-message>New person #{person.firstName} #{person.lastName} created</framework:created-message>
    <framework:deleted-message>Person #{person.firstName} #{person.lastName} deleted</framework:deleted-message>
    <framework:updated-message>Person #{person.firstName} #{person.lastName} updated</framework:updated-message>
</framework:entity-home>

<component name="newPerson"
           class="eg.Person">
    <property name="nationality">#{country}</property>
</component>

```

Or extension:

```

@Name("personHome")
public class PersonHome extends EntityHome<Person> {

    @In Country country;

    @Factory("person")
    public Person initPerson() { return getInstance(); }

    protected Person createInstance() {

```

```

        return new Person(country);
    }

    protected String getCreatedMessage() { return "New person #{person.firstName} #{person.lastName} created"; }
    protected String getUpdatedMessage() { return "Person #{person.firstName} #{person.lastName} updated"; }
    protected String getDeletedMessage() { return "Person #{person.firstName} #{person.lastName} deleted"; }
}

```

But the best way to specify the messages is to put them in a resource bundle known to Seam (the bundle named `messages`, by default).

```

Person_created=New person #{person.firstName} #{person.lastName} created
Person_deleted=Person #{person.firstName} #{person.lastName} deleted
Person_updated=Person #{person.firstName} #{person.lastName} updated

```

This enables internationalization, and keeps your code and configuration clean of presentation concerns.

The final step is to add validation functionality to the page, using `<s:validateAll>` and `<s:decorate>`, but I'll leave that for you to figure out.

11.3. Query objects

If we need a list of all `Person` instance in the database, we can use a `Query` object. For example:

```

<framework:entity-query name="people"
    ejbql="select p from Person p"/>

```

We can use it from a JSF page:

```

<h1>List of people</h1>
<h:dataTable value="#{people.resultList}" var="person">
    <h:column>
        <s:link view="/editPerson.jsp" value="#{person.firstName} #{person.lastName}">
            <f:param name="personId" value="#{person.id}"/>
        </s:link>
    </h:column>
</h:dataTable>

```

We probably need to support pagination:

```

<framework:entity-query name="people"
    ejbql="select p from Person p"
    order="lastName"
    max-results="20"/>

```

We'll use a page parameter to determine the page to display:

```

<pages>
    <page view-id="/searchPerson.jsp">
        <param name="firstResult" value="#{people.firstResult}"/>
    </page>
</pages>

```

The JSF code for a pagination control is a bit verbose, but manageable:

```

<h1>Search for people</h1>
<h:dataTable value="#{people.resultList}" var="person">
    <h:column>
        <s:link view="/editPerson.jsp" value="#{person.firstName} #{person.lastName}">

```

```

        <f:param name="personId" value="#{person.id}"/>
    </s:link>
</h:column>
</h:dataTable>

<s:link view="/search.xhtml" rendered="#{people.previousExists}" value="First Page">
    <f:param name="firstResult" value="0"/>
</s:link>

<s:link view="/search.xhtml" rendered="#{people.previousExists}" value="Previous Page">
    <f:param name="firstResult" value="#{people.previousFirstResult}"/>
</s:link>

<s:link view="/search.xhtml" rendered="#{people.nextExists}" value="Next Page">
    <f:param name="firstResult" value="#{people.nextFirstResult}"/>
</s:link>

<s:link view="/search.xhtml" rendered="#{people.nextExists}" value="Last Page">
    <f:param name="firstResult" value="#{people.lastFirstResult}"/>
</s:link>

```

Real search screens let the user enter a bunch of optional search criteria to narrow the list of results returned. The Query object lets you specify optional "restrictions" to support this important usecase:

```

<component name="examplePerson" class="Person"/>

<framework:entity-query name="people"
    ejbql="select p from Person p"
    order="lastName"
    max-results="20">
    <framework:restrictions>
        <value>lower(firstName) like lower( concat("#{examplePerson.firstName},'%') )</value>
        <value>lower(lastName) like lower( concat("#{examplePerson.lastName},'%') )</value>
    </framework:restrictions>
</framework:entity-query>

```

Notice the use of an "example" object.

```

<h1>Search for people</h1>
<h:form>
    <div>First name: <h:inputText value="#{examplePerson.firstName}"/></div>
    <div>Last name: <h:inputText value="#{examplePerson.lastName}"/></div>
    <div><h:commandButton value="Search" action="/search.jsp"/></div>
</h:form>

<h:dataTable value="#{people.resultList}" var="person">
    <h:column>
        <s:link view="/editPerson.jsp" value="#{person.firstName} #{person.lastName}">
            <f:param name="personId" value="#{person.id}"/>
        </s:link>
    </h:column>
</h:dataTable>

```

To refresh the query when the underlying entities change we observe the `org.jboss.seam.afterTransactionSuccess` event:

```

<event type="org.jboss.seam.afterTransactionSuccess">
    <action execute="#{people.refresh}" />
</event>

```

Or, to just refresh the query when the person entity is persisted, updated or removed through `PersonHome`:

```

<event type="org.jboss.seam.afterTransactionSuccess.Person">
    <action execute="#{people.refresh}" />
</event>

```

The examples in this section have all shown reuse by configuration. However, reuse by extension is equally possible for Query objects.

11.4. Controller objects

A totally optional part of the Seam Application Framework is the class `Controller` and its subclasses `EntityController`, `HibernateEntityController` and `BusinessProcessController`. These classes provide nothing more than some convenience methods for access to commonly used built-in components and methods of built-in components. They help save a few keystrokes (characters can add up!) and provide a great launchpad for new users to explore the rich functionality built in to Seam.

For example, here is what `RegisterAction` from the Seam registration example would look like:

```
@Stateless
@Name("register")
public class RegisterAction extends EntityController implements Register
{
    @In private User user;

    public String register()
    {
        List existing = createQuery("select u.username from User u where u.username=:username")
            .setParameter("username", user.getUsername())
            .getResultList();

        if ( existing.size()==0 )
        {
            persist(user);
            info("Registered new user #{user.username}");
            return "/registered.jspx";
        }
        else
        {
            addFacesMessage("User #{user.username} already exists");
            return null;
        }
    }
}
```

As you can see, its not an earthshattering improvement...

Chapter 12. Seam and JBoss Rules

Seam makes it easy to call JBoss Rules (Drools) rulebases from Seam components or jBPM process definitions.

12.1. Installing rules

The first step is to make an instance of `org.drools.RuleBase` available in a Seam context variable. For testing purposes, Seam provides a built-in component that compiles a static set of rules from the classpath. You can install this component via `components.xml`:

```
<drools:rule-base name="policyPricingRules">
  <drools:rule-files>
    <value>policyPricingRules</value>
  </drools:rule-files>
</drools:rule-base>
```

This component compiles rules from a set of `.dr1` files and caches an instance of `org.drools.RuleBase` in the Seam `APPLICATION` context. Note that it is quite likely that you will need to install multiple rule bases in a rule-driven application.

If you want to use a Drools DSL, you also need to specify the DSL definition:

```
<drools:rule-base name="policyPricingRules" dsl-file="policyPricing.dsl">
  <drools:rule-files>
    <value>policyPricingRules</value>
  </drools:rule-files>
</drools:rule-base>
```

In most rules-driven applications, rules need to be dynamically deployable, so a production application will want to use a Drools `RuleAgent` to manage the `RuleBase`. The `RuleAgent` can connect to a Drools rule server (BRMS) or hot deploy rules packages from a local file repository. The `RulesAgent`-managed `RuleBase` is also configurable in `components.xml`:

```
<drools:rule-agent name="insuranceRules"
  configurationFile="/WEB-INF/deployedrules.properties" />
```

The properties file contains properties specific to the `RulesAgent`. Here is an example configuration file from the Drools example distribution.

```
newInstance=true
url=http://localhost:8080/drools-jbrms/org.drools.brms.JBRMS/package/org.acme.insurance/fmeyer
localCacheDir=/Users/fernandomeyer/projects/jbossrules/drools-examples/drools-examples-brms/cache
poll=30
name=insuranceconfig
```

It is also possible to configure the options on the component directly, bypassing the configuration file.

```
<drools:rule-agent name="insuranceRules"
  url="http://localhost:8080/drools-jbrms/org.drools.brms.JBRMS/package/org.acme.insurance"
  local-cache-dir="/Users/fernandomeyer/projects/jbossrules/drools-examples/drools-examples-brms/cache"
  poll="30"
  configuration-name="insuranceconfig" />
```

Next, we need to make an instance of `org.drools.WorkingMemory` available to each conversation. (Each `WorkingMemory` accumulates facts relating to the current conversation.)

```
<drools:managed-working-memory name="policyPricingWorkingMemory" auto-create="true" rule-base="#{policyPricingRuleBase}" />
```

Notice that we gave the `policyPricingWorkingMemory` a reference back to our rule base via the `ruleBase` configuration property.

12.2. Using rules from a Seam component

We can now inject our `WorkingMemory` into any Seam component, assert facts, and fire rules:

```
@In WorkingMemory policyPricingWorkingMemory;

@In Policy policy;
@In Customer customer;

public void pricePolicy() throws FactException
{
    policyPricingWorkingMemory.assertObject(policy);
    policyPricingWorkingMemory.assertObject(customer);
    policyPricingWorkingMemory.fireAllRules();
}
```

12.3. Using rules from a jBPM process definition

You can even allow a rule base to act as a jBPM action handler, decision handler, or assignment handler—in either a pageflow or business process definition.

```
<decision name="approval">

    <handler class="org.jboss.seam.drools.DroolsDecisionHandler">
        <workingMemoryName>orderApprovalRulesWorkingMemory</workingMemoryName>
        <assertObjects>
            <element>#{customer}</element>
            <element>#{order}</element>
            <element>#{order.lineItems}</element>
        </assertObjects>
    </handler>

    <transition name="approved" to="ship">
        <action class="org.jboss.seam.drools.DroolsActionHandler">
            <workingMemoryName>shippingRulesWorkingMemory</workingMemoryName>
            <assertObjects>
                <element>#{customer}</element>
                <element>#{order}</element>
                <element>#{order.lineItems}</element>
            </assertObjects>
        </action>
    </transition>

    <transition name="rejected" to="cancelled"/>
</decision>
```

The `<assertObjects>` element specifies EL expressions that return an object or collection of objects to be asserted as facts into the `WorkingMemory`.

There is also support for using Drools for jBPM task assignments:

```
<task-node name="review">
    <task name="review" description="Review Order">
        <assignment handler="org.jboss.seam.drools.DroolsAssignmentHandler">
```

```

        <workingMemoryName>orderApprovalRulesWorkingMemory</workingMemoryName>
        <assertObjects>
            <element>#{actor}</element>
            <element>#{customer}</element>
            <element>#{order}</element>
            <element>#{order.lineItems}</element>
        </assertObjects>
    </assignment>
</task>
<transition name="rejected" to="cancelled"/>
<transition name="approved" to="approved"/>
</task-node>

```

Certain objects are available to the rules as Drools globals, namely the jBPM Assignable, as assignable and a Seam Decision object, as decision. Rules which handle decisions should call `decision.setOutcome("result")` to determine the result of the decision. Rules which perform assignments should set the actor id using the Assignable.

```

package org.jboss.seam.examples.shop

import org.jboss.seam.drools.Decision

global Decision decision

rule "Approve Order For Loyal Customer"
    when
        Customer( loyaltyStatus == "GOLD" )
        Order( totalAmount <= 10000 )
    then
        decision.setOutcome("approved");
    end

```

```

package org.jboss.seam.examples.shop

import org.jbpm.taskmgmt.exe.Assignable

global Assignable assignable

rule "Assign Review For Small Order"
    when
        Order( totalAmount <= 100 )
    then
        assignable.setPooledActors( new String[] {"reviewers"} );
    end

```

Chapter 13. Security

The Seam Security API is an optional Seam feature that provides authentication and authorization features for securing both domain and page resources within your Seam project.

13.1. Overview

Seam Security provides two different modes of operation:

- *simplified mode* - this mode supports authentication services and simple role-based security checks.
- *advanced mode* - this mode supports all the same features as the simplified mode, plus it offers rule-based security checks using JBoss Rules.

13.1.1. Which mode is right for my application?

That all depends on the requirements of your application. If you have minimal security requirements, for example if you only wish to restrict certain pages and actions to users who are logged in, or who belong to a certain role, then the simplified mode will probably be sufficient. The advantages of this is a more simplified configuration, significantly less libraries to include, and a smaller memory footprint.

If on the other hand, your application requires security checks based on contextual state or complex business rules, then you will require the features provided by the advanced mode.

13.2. Requirements

If using the advanced mode features of Seam Security, the following jar files are required to be configured as modules in `application.xml`. If you are using Seam Security in simplified mode, these are *not* required:

- `drools-compiler.jar`
- `drools-core.jar`
- `janino.jar`
- `antlr-runtime.jar`
- `mvel14.jar`

For web-based security, `jboss-seam-ui.jar` must also be included in the application's war file.

13.3. Disabling Security

In some situations it may be necessary to disable Seam Security, for example during unit tests. This can be done by calling the static method `Identity.setSecurityEnabled(false)` to disable security checks. Doing this prevents any security checks being performed for the following:

- Entity Security
- Hibernate Security Interceptor
- Seam Security Interceptor
- Page restrictions

13.4. Authentication

The authentication features provided by Seam Security are built upon JAAS (Java Authentication and Authorization Service), and as such provide a robust and highly configurable API for handling user authentication. However, for less complex authentication requirements Seam offers a much more simplified method of authentication that hides the complexity of JAAS.

13.4.1. Configuration

The simplified authentication method uses a built-in JAAS login module, `SeamLoginModule`, which delegates authentication to one of your own Seam components. This login module is already configured inside Seam as part of a default application policy and as such does not require any additional configuration files. It allows you to write an authentication method using the entity classes that are provided by your own application. Configuring this simplified form of authentication requires the identity component to be configured in `components.xml`:

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:security="http://jboss.com/products/seam/security"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation=
    "http://jboss.com/products/seam/components http://jboss.com/products/seam/components-2.0.xsd
    http://jboss.com/products/seam/security http://jboss.com/products/seam/security-2.0.xsd"

  <security:identity authenticate-method="#{authenticator.authenticate}"/>

</components>
```

If you wish to use the advanced security features such as rule-based permission checks, all you need to do is include the Drools (JBoss Rules) jars in your classpath, and add some additional configuration, described later.

The EL expression `#{authenticator.authenticate}` is a method binding indicating that the `authenticate` method of the `authenticator` component will be used to authenticate the user.

13.4.2. Writing an authentication method

The `authenticate-method` property specified for `identity` in `components.xml` specifies which method will be used by `SeamLoginModule` to authenticate users. This method takes no parameters, and is expected to return a boolean indicating whether authentication is successful or not. The user's username and password can be obtained from `Identity.instance().getUsername()` and `Identity.instance().getPassword()`, respectively. Any roles that the user is a member of should be assigned using `Identity.instance().addRole()`. Here's a complete example of an authentication method inside a JavaBean component:

```
@Name("authenticator")
public class Authenticator {
    @In EntityManager entityManager;
```

```
public boolean authenticate() {
    try
    {
        User user = (User) entityManager.createQuery(
            "from User where username = :username and password = :password")
            .setParameter("username", Identity.instance().getUsername())
            .setParameter("password", Identity.instance().getPassword())
            .getSingleResult();


        if (user.getRoles() != null)
        {
            for (UserRole mr : user.getRoles())
                Identity.instance().addRole(mr.getName());
        }

        return true;
    }
    catch (NoResultException ex)
    {
        return false;
    }
}
```

In the above example, both `User` and `UserRole` are application-specific entity beans. The `roles` parameter is populated with the roles that the user is a member of, which should be added to the `Set` as literal string values, e.g. "admin", "user". In this case, if the user record is not found and a `NoResultException` thrown, the authentication method returns `false` to indicate the authentication failed.

Identity.addRole()

The `Identity.addRole()` method behaves differently depending on whether the current session is authenticated or not. If the session is not authenticated, then `addRole()` should *only* be called during the authentication process. When called here, the role name is placed into a temporary list of pre-authenticated roles. Once authentication is successful, the pre-authenticated roles then become "real" roles, and calling `Identity.hasRole()` for those roles will then return `true`. The following sequence diagram represents the list of pre-authenticated roles as a first class object to show more clearly how it fits in to the authentication process.



Special Considerations

When writing an authenticator method, it is important that it is kept minimal and free from any side-effects. This is because there is no guarantee as to how many times the authenticator method will be called by the security API, and as such it may be invoked multiple times during a single request. Because of this, any special code that should execute upon a successful or failed authentication should be written by implementing an event observer. See the section on Security Events further down in this chapter for more information about which events are raised by Seam Security.

To give an example, let's say that upon a successful login that some user statistics must be updated. We would do this by writing an event observer for the `org.jboss.seam.security.loginSuccessful` event, like this:

```

@In UserStats userStats;

@Observer("org.jboss.seam.security.loginSuccessful")
public void updateUserStats()
{
    userStats.setLastLoginDate(new Date());
    userStats.incrementLoginCount();
}

```

13.4.3. Writing a login form

The Identity component provides both username and password properties, catering for the most common authentication scenario. These properties can be bound directly to the username and password fields on a login form. Once these properties are set, calling the `identity.login()` method will authenticate the user using the provided credentials. Here's an example of a simple login form:

```

<div>
    <h:outputLabel for="name" value="Username"/>

```

```

    <h:inputText id="name" value="#{identity.username}"/>
</div>

<div>
    <h:outputLabel for="password" value="Password"/>
    <h:inputSecret id="password" value="#{identity.password}"/>
</div>

<div>
    <h:commandButton value="Login" action="#{identity.login}"/>
</div>

```

Similarly, logging out the user is done by calling `#{identity.logout}`. Calling this action will clear the security state of the currently authenticated user.

13.4.4. Simplified Configuration - Summary

So to sum up, there are the three easy steps to configure authentication:

- Configure an authentication method in `components.xml`.
- Write an authentication method.
- Write a login form so that the user can authenticate.

13.4.5. Handling Security Exceptions

To prevent users from receiving the default error page in response to a security error, it's recommended that `pages.xml` is configured to redirect security errors to a more "pretty" page. The two main types of exceptions thrown by the security API are:

- `NotLoggedInException` - This exception is thrown if the user attempts to access a restricted action or page when they are not logged in.
- `AuthorizationException` - This exception is only thrown if the user is already logged in, and they have attempted to access a restricted action or page for which they do not have the necessary privileges.

In the case of a `NotLoggedInException`, it is recommended that the user is redirected to either a login or registration page so that they can log in. For an `AuthorizationException`, it may be useful to redirect the user to an error page. Here's an example of a `pages.xml` file that redirects both of these security exceptions:

```

<pages>

    ...

    <exception class="org.jboss.seam.security.NotLoggedInException">
        <redirect view-id="/login.xhtml">
            <message>You must be logged in to perform this action</message>
        </redirect>
    </exception>

    <exception class="org.jboss.seam.security.AuthorizationException">
        <end-conversation/>
        <redirect view-id="/security_error.xhtml">
            <message>You do not have the necessary security privileges to perform this action.</message>
        </redirect>
    </exception>

```

```
</pages>
```

Most web applications require even more sophisticated handling of login redirection, so Seam includes some special functionality for handling this problem.

13.4.6. Login Redirection

You can ask Seam to redirect the user to a login screen when an unauthenticated user tries to access a particular view (or wildcarded view id) as follows:

```
<pages login-view-id="/login.xhtml">
    <page view-id="/members/*" login-required="true"/>
    ...
</pages>
```

(This is less of a blunt instrument than the exception handler shown above, but should probably be used in conjunction with it.)

After the user logs in, we want to automatically send them back where they came from, so they can retry the action that required logging in. If you add the following event listeners to `components.xml`, attempts to access a restricted view while not logged in will be remembered, so that upon the user successfully logging in they will be redirected to the originally requested view, with any page parameters that existed in the original request.

```
<event type="org.jboss.seam.security.notLoggedIn">
    <action execute="#{redirect.captureCurrentView}"/>
</event>

<event type="org.jboss.seam.security.postAuthenticate">
    <action execute="#{redirect.returnToCapturedView}"/>
</event>
```

Note that login redirection is implemented as a conversation-scoped mechanism, so don't end the conversation in your `authenticate()` method.

13.4.7. HTTP Authentication

Although not recommended for use unless absolutely necessary, Seam provides means for authenticating using either HTTP Basic or HTTP Digest (RFC 2617) methods. To use either form of authentication, the `authentication-filter` component must be enabled in `components.xml`:

```
<web:authentication-filter url-pattern="*.seam" auth-type="basic"/>
```

To enable the filter for basic authentication, set `auth-type` to `basic`, or for digest authentication, set it to `digest`. If using digest authentication, the `key` and `realm` must also be set:

```
<web:authentication-filter url-pattern="*.seam" auth-type="digest" key="AA3JK34aSDlkj" realm="My App"/>
```

The `key` can be any String value. The `realm` is the name of the authentication realm that is presented to the user when they authenticate.

Writing a Digest Authenticator

If using digest authentication, your authenticator class should extend the abstract class `org.jboss.seam.security.digest.DigestAuthenticator`, and use the `validatePassword()` method to validate the user's plain text password against the digest request. Here is an example:

```
public boolean authenticate()
{
    try
    {
        User user = (User) entityManager.createQuery(
            "from User where username = :username")
            .setParameter("username", identity.getUsername())
            .getSingleResult();

        return validatePassword(user.getPassword());
    }
    catch (NoResultException ex)
    {
        return false;
    }
}
```

13.4.8. Advanced Authentication Features

This section explores some of the advanced features provided by the security API for addressing more complex security requirements.

Using your container's JAAS configuration

If you would rather not use the simplified JAAS configuration provided by the Seam Security API, you may instead delegate to the default system JAAS configuration by providing a `jaasConfigName` property in `components.xml`. For example, if you are using JBoss AS and wish to use the other policy (which uses the `UserRolesLoginModule` login module provided by JBoss AS), then the entry in `components.xml` would look like this:

```
<security:identity authenticate-method="#{authenticator.authenticate}"
    jaas-config-name="other"/>
```

13.5. Error Messages

The security API produces a number of default faces messages for various security-related events. The following table lists the message keys that can be used to override these messages by specifying them in a `message.properties` resource file. To suppress the message, just put the key with an empty value in the resource file.

Table 13.1. Security Message Keys

Message Key	Description
<code>org.jboss.seam.security.loginSuccessful</code>	This message is produced when a user successfully logs in via the security API.

Message Key	Description
org.jboss.seam.security.loginFailed	This message is produced when the login process fails, either because the user provided an incorrect username or password, or because authentication failed in some other way.
org.jboss.seam.security.NotLoggedIn	This message is produced when a user attempts to perform an action or access a page that requires a security check, and the user is not currently authenticated.

13.6. Authorization

There are a number of authorization features provided by the Seam Security API for securing access to components, component methods, and pages. This section describes each of these. An important thing to note is that if you wish to use any of the advanced features (such as rule-based permissions) then your `components.xml` must be configured to support this - see the Configuration section above.

13.6.1. Core concepts

Each of the authorization mechanisms provided by the Seam Security API are built upon the concept of a user being granted roles and/or permissions. A role is a *group*, or *type*, of user that may have been granted certain privileges for performing one or more specific actions within an application. A permission on the other hand is a privilege (sometimes once-off) for performing a single, specific action. It is entirely possible to build an application using nothing but permissions, however roles offer a higher level of convenience when granting privileges to groups of users.

Roles are simple, consisting of only a name such as "admin", "user", "customer", etc. Permissions consist of both a name and an action, and are represented within this documentation in the form `name:action`, for example `customer:delete`, or `customer:insert`.

13.6.2. Securing components

Let's start by examining the simplest form of authorization, component security, starting with the `@Restrict` annotation.

The `@Restrict` annotation

Seam components may be secured either at the method or the class level, using the `@Restrict` annotation. If both a method and it's declaring class are annotated with `@Restrict`, the method restriction will take precedence (and the class restriction will not apply). If a method invocation fails a security check, then an exception will be thrown as per the contract for `Identity.checkRestriction()` (see Inline Restrictions). A `@Restrict` on just the component class itself is equivalent to adding `@Restrict` to each of its methods.

An empty `@Restrict` implies a permission check of `componentName:methodName`. Take for example the following component method:

```
@Name("account")
public class AccountAction {
    @Restrict public void delete() {
        ...
    }
}
```

```
}

```

In this example, the implied permission required to call the `delete()` method is `account:delete`. The equivalent of this would be to write `@Restrict("#{s:hasPermission('account','delete',null)}")`. Now let's look at another example:

```
@Restrict @Name("account")
public class AccountAction {
    public void insert() {
        ...
    }
    @Restrict("#{s:hasRole('admin')}")
    public void delete() {
        ...
    }
}
```

This time, the component class itself is annotated with `@Restrict`. This means that any methods without an overriding `@Restrict` annotation require an implicit permission check. In the case of this example, the `insert()` method requires a permission of `account:insert`, while the `delete()` method requires that the user is a member of the `admin` role.

Before we go any further, let's address the `#{s:hasRole()}` expression seen in the above example. Both `s:hasRole` and `s:hasPermission` are EL functions, which delegate to the correspondingly named methods of the `Identity` class. These functions can be used within any EL expression throughout the entirety of the security API.

Being an EL expression, the value of the `@Restrict` annotation may reference any objects that exist within a Seam context. This is extremely useful when performing permission checks for a specific object instance. Look at this example:

```
@Name("account")
public class AccountAction {
    @In Account selectedAccount;
    @Restrict("#{s:hasPermission('account','modify',selectedAccount)}")
    public void modify() {
        selectedAccount.modify();
    }
}
```

The interesting thing to note from this example is the reference to `selectedAccount` seen within the `hasPermission()` function call. The value of this variable will be looked up from within the Seam context, and passed to the `hasPermission()` method in `Identity`, which in this case can then determine if the user has the required permission for modifying the specified `Account` object.

Inline restrictions

Sometimes it might be desirable to perform a security check in code, without using the `@Restrict` annotation. In this situation, simply use `Identity.checkRestriction()` to evaluate a security expression, like this:

```
public void deleteCustomer() {
    Identity.instance().checkRestriction("#{s:hasPermission('customer','delete',selectedCustomer)}");
}
```

If the expression specified doesn't evaluate to `true`, either

- if the user is not logged in, a `NotLoggedInException` exception is thrown or

- if the user is logged in, an `AuthorizationException` exception is thrown.

It is also possible to call the `hasRole()` and `hasPermission()` methods directly from Java code:

```
if (!Identity.instance().hasRole("admin"))
    throw new AuthorizationException("Must be admin to perform this action");

if (!Identity.instance().hasPermission("customer", "create", null))
    throw new AuthorizationException("You may not create new customers");
```

13.6.3. Security in the user interface

One indication of a well designed user interface is that the user is not presented with options for which they don't have the necessary privileges to use. Seam Security allows conditional rendering of either 1) sections of a page or 2) individual controls, based upon the privileges of the user, using the very same EL expressions that are used for component security.

Let's take a look at some examples of interface security. First of all, let's pretend that we have a login form that should only be rendered if the user is not already logged in. Using the `identity.isLoggedIn()` property, we can write this:

```
<h:form class="loginForm" rendered="#{not identity.loggedIn}">
```

If the user isn't logged in, then the login form will be rendered - very straight forward so far. Now let's pretend there is a menu on the page that contains some actions which should only be accessible to users in the manager role. Here's one way that these could be written:

```
<h:outputLink action="#{reports.listManagerReports}" rendered="#{s:hasRole('manager')}">
    Manager Reports
</h:outputLink>
```

This is also quite straight forward. If the user is not a member of the manager role, then the `outputLink` will not be rendered. The `rendered` attribute can generally be used on the control itself, or on a surrounding `<s:div>` or `<s:span>` control.

Now for something more complex. Let's say you have a `h:dataTable` control on a page listing records for which you may or may not wish to render action links depending on the user's privileges. The `s:hasPermission` EL function allows us to pass in an object parameter which can be used to determine whether the user has the requested permission for that object or not. Here's how a `dataTable` with secured links might look:

```
<h:dataTable value="#{clients}" var="cl">
    <h:column>
        <f:facet name="header">Name</f:facet>
        #{cl.name}
    </h:column>
    <h:column>
        <f:facet name="header">City</f:facet>
        #{cl.city}
    </h:column>
    <h:column>
        <f:facet name="header">Action</f:facet>
        <s:link value="Modify Client" action="#{clientAction.modify}"
            rendered="#{s:hasPermission('client','modify',cl)}/>
        <s:link value="Delete Client" action="#{clientAction.delete}"
            rendered="#{s:hasPermission('client','delete',cl)}/>
    </h:column>
</h:dataTable>
```

13.6.4. Securing pages

Page security requires that the application is using a `pages.xml` file, however is extremely simple to configure. Simply include a `<restrict/>` element within the page elements that you wish to secure. If no explicit restriction is specified by the `restrict` element, an implied permission of `/viewId.xhtml:render` will be checked when the page is accessed via a non-faces (GET) request, and a permission of `/viewId.xhtml:restore` will be required when any JSF postback (form submission) originates from the page. Otherwise, the specified restriction will be evaluated as a standard security expression. Here's a couple of examples:

```
<page view-id="/settings.xhtml">
  <restrict/>
</page>
```

This page has an implied permission of `/settings.xhtml:render` required for non-faces requests and an implied permission of `/settings.xhtml:restore` for faces requests.

```
<page view-id="/reports.xhtml">
  <restrict>#{s:hasRole('admin')}}</restrict>
</page>
```

Both faces and non-faces requests to this page require that the user is a member of the admin role.

13.6.5. Securing Entities

Seam security also makes it possible to apply security restrictions to read, insert, update and delete actions for entities.

To secure all actions for an entity class, add a `@Restrict` annotation on the class itself:

```
@Entity
@Name("customer")
@Restrict
public class Customer {
    ...
}
```

If no expression is specified in the `@Restrict` annotation, the default security check that is performed is a permission check of `entityName:action`, where `entityName` is the Seam component name of the entity (or the fully-qualified class name if no `@Name` is specified), and the action is either read, insert, update or delete.

It is also possible to only restrict certain actions, by placing a `@Restrict` annotation on the relevant entity life-cycle method (annotated as follows):

- `@PostLoad` - Called after an entity instance is loaded from the database. Use this method to configure a read permission.
- `@PrePersist` - Called before a new instance of the entity is inserted. Use this method to configure an insert permission.
- `@PreUpdate` - Called before an entity is updated. Use this method to configure an update permission.
- `@PreRemove` - Called before an entity is deleted. Use this method to configure a delete permission.

Here's an example of how an entity would be configured to perform a security check for any insert operations. Please note that the method is not required to do anything, the only important thing in regard to security is how

it is annotated:

```
@PrePersist @Restrict
public void prePersist() {}
```

And here's an example of an entity permission rule that checks if the authenticated user is allowed to insert a new MemberBlog record (from the seam space example). The entity for which the security check is being made is automatically inserted into the working memory (in this case MemberBlog):

```
rule InsertMemberBlog
  no-loop
  activation-group "permissions"
when
  check: PermissionCheck(name == "memberBlog", action == "insert", granted == false)
  Principal(principalName : name)
  MemberBlog(member : member -> (member.getUsername().equals(principalName)))
then
  check.grant();
end;
```

This rule will grant the permission memberBlog:insert if the currently authenticated user (indicated by the Principal fact) has the same name as the member for which the blog entry is being created. The "principalName : name" structure that can be seen in the Principal fact (and other places) is a variable binding - it binds the name property of the Principal to a variable called principalName. Variable bindings allow the value to be referred to in other places, such as the following line which compares the member's username to the Principal name. For more details, please refer to the JBoss Rules documentation.

Finally, we need to install a listener class that integrates Seam security with your JPA provider.

Entity security with JPA

Security checks for EJB3 entity beans are performed with an EntityListener. You can install this listener by using the following META-INF/orm.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://java.sun.com/xml/ns/persistence/orm"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/persistence/orm http://java.sun.com/xml/ns/persistence/orm-1.0.xsd"
  version="1.0">

  <persistence-unit-metadata>
    <persistence-unit-defaults>
      <entity-listeners>
        <entity-listener class="org.jboss.seam.security.EntitySecurityListener"/>
      </entity-listeners>
    </persistence-unit-defaults>
  </persistence-unit-metadata>

</entity-mappings>
```

Entity security with Hibernate

If you are using a Hibernate SessionFactory configured via Seam, you don't need to do anything special to use entity security.

13.7. Writing Security Rules

Up to this point there has been a lot of mention of permissions, but no information about how permissions are actually defined or granted. This section completes the picture, by explaining how permission checks are processed, and how to implement permission checks for a Seam application.

13.7.1. Permissions Overview

So how does the security API know whether a user has the `customer:modify` permission for a specific customer? Seam Security provides quite a novel method for determining user permissions, based on JBoss Rules. A couple of the advantages of using a rule engine are 1) a centralized location for the business logic that is behind each user permission, and 2) speed - JBoss Rules uses very efficient algorithms for evaluating large numbers of complex rules involving multiple conditions.

13.7.2. Configuring a rules file

Seam Security expects to find a `RuleBase` component called `securityRules` which it uses to evaluate permission checks. This is configured in `components.xml` as follows:

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:security="http://jboss.com/products/seam/security"
  xmlns:drools="http://jboss.com/products/seam/drools"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation=
    "http://jboss.com/products/seam/core http://jboss.com/products/seam/core-2.0.xsd
    http://jboss.com/products/seam/components http://jboss.com/products/seam/components-2.0.xsd
    http://jboss.com/products/seam/drools http://jboss.com/products/seam/drools-2.0.xsd
    http://jboss.com/products/seam/security http://jboss.com/products/seam/security-2.0.xsd"
  >
  <drools:rule-base name="securityRules">
    <drools:rule-files>
      <value>/META-INF/security.drl</value>
    </drools:rule-files>
  </drools:rule-base>
</components>
```

Once the `RuleBase` component is configured, it's time to write the security rules.

13.7.3. Creating a security rules file

For this step you need to create a file called `security.drl` in the `/META-INF` directory of your application's jar file. In actual fact this file can be called anything you want, and exist in any location as long as it is configured appropriately in `components.xml`.

So what should the security rules file contain? At this stage it might be a good idea to at least skim through the JBoss Rules documentation, however to get started here's an extremely simple example:

```
package MyApplicationPermissions;

import org.jboss.seam.security.PermissionCheck;
import org.jboss.seam.security.Role;

rule CanUserDeleteCustomers
when
  c: PermissionCheck(name == "customer", action == "delete")
  Role(name == "admin")
then
  c.grant();
end;
```

Let's break this down. The first thing we see is the package declaration. A package in JBoss Rules is essentially a collection of rules. The package name can be anything you want - it doesn't relate to anything else outside the scope of the rule base.

The next thing we can notice is a couple of import statements for the `PermissionCheck` and `Role` classes. These imports inform the rules engine that we'll be referencing these classes within our rules.

Finally we have the code for the rule. Each rule within a package should be given a unique name (usually describing the purpose of the rule). In this case our rule is called `CanUserDeleteCustomers` and will be used to check whether a user is allowed to delete a customer record.

Looking at the body of the rule definition we can notice two distinct sections. Rules have what is known as a left hand side (LHS) and a right hand side (RHS). The LHS consists of the conditional part of the rule, i.e. a list of conditions which must be satisfied for the rule to fire. The LHS is represented by the `when` section. The RHS is the consequence, or action section of the rule that will only be fired if all of the conditions in the LHS are met. The RHS is represented by the `then` section. The end of the rule is denoted by the `end;` line.

If we look at the LHS of the rule, we see two conditions listed there. Let's examine the first condition:

```
c: PermissionCheck(name == "customer", action == "delete")
```

In plain english, this condition is stating that there must exist a `PermissionCheck` object with a `name` property equal to `"customer"`, and an `action` property equal to `"delete"` within the working memory. What is the working memory? It is a session-scoped object that contains the contextual information that is required by the rules engine to make a decision about a permission check. Each time the `hasPermission()` method is called, a temporary `PermissionCheck` object, or *Fact*, is inserted into the working memory. This `PermissionCheck` corresponds exactly to the permission that is being checked, so for example if you call `hasPermission("account", "create", null)` then a `PermissionCheck` object with a `name` equal to `"account"` and `action` equal to `"create"` will be inserted into the working memory for the duration of the permission check.

So what else is in the working memory? Besides the short-lived temporary facts inserted during a permission check, there are some longer-lived objects in the working memory that stay there for the entire duration of a user being authenticated. These include any `java.security.Principal` objects that are created as part of the authentication process, plus a `org.jboss.seam.security.Role` object for each of the roles that the user is a member of. It is also possible to insert additional long-lived facts into the working memory by calling `((RuleBasedIdentity) RuleBasedIdentity.instance()).getSecurityContext().insert()`, passing the object as a parameter.

Getting back to our simple example, we can also notice that the first line of our LHS is prefixed with `c:`. This is a variable binding, and is used to refer back to the object that is matched by the condition. Moving onto the second line of our LHS, we see this:

```
Role(name == "admin")
```

This condition simply states that there must be a `Role` object with a `name` of `"admin"` within the working memory. As mentioned, user roles are inserted into the working memory as long-lived facts. So, putting both conditions together, this rule is essentially saying "I will fire if you are checking for the `customer:delete` permission and the user is a member of the `admin` role".

So what is the consequence of the rule firing? Let's take a look at the RHS of the rule:

```
c.grant()
```

The RHS consists of Java code, and in this case is invoking the `grant()` method of the `c` object, which as

already mentioned is a variable binding for the `PermissionCheck` object. Besides the name and action properties of the `PermissionCheck` object, there is also a `granted` property which is initially set to `false`. Calling `grant()` on a `PermissionCheck` sets the `granted` property to `true`, which means that the permission check was successful, allowing the user to carry out whatever action the permission check was intended for.

Wildcard permission checks

It is possible to implement a wildcard permission check (which allows all actions for a given permission name), by omitting the action constraint for the `PermissionCheck` in your rule, like this:

```
rule CanDoAnythingToCustomersIfYouAreAnAdmin
when
  c: PermissionCheck(name == "customer")
  Role(name == "admin")
then
  c.grant();
end;
```

This rule allows users with the `admin` role to perform *any* action for any customer permission check.

13.8. SSL Security

Seam includes basic support for serving sensitive pages via the HTTPS protocol. This is easily configured by specifying a scheme for the page in `pages.xml`. The following example shows how the view `/login.xhtml` is configured to use HTTPS:

```
<page view-id="/login.xhtml" scheme="https"/>
```

This configuration is automatically extended to both `s:link` and `s:button` JSF controls, which (when specifying the view) will also render the link using the correct protocol. Based on the previous example, the following link will use the HTTPS protocol because `/login.xhtml` is configured to use it:

```
<s:link view="/login.xhtml" value="Login"/>
```

Browsing directly to a view when using the *incorrect* protocol will cause a redirect to the same view using the *correct* protocol. For example, browsing to a page that has `scheme="https"` using HTTP will cause a redirect to the same page using HTTPS.

It is also possible to configure a *default scheme* for all pages. This is useful if you wish to use HTTPS for a only few pages. If no default scheme is specified then the normal behavior is to continue use the current scheme. So once the user accessed a page that required HTTPS, then HTTPS would continue to be used after the user navigated away to other non-HTTPS pages. (While this is good for security, it is not so great for performance!). To define HTTP as the default scheme, add this line to `pages.xml`:

```
<page view-id="*" scheme="http" />
```

Of course, if *none* of the pages in your application use HTTPS then it is not required to specify a default scheme.

You may configure Seam to automatically invalidate the current HTTP session each time the scheme changes. Just add this line to `components.xml`:

```
<core:servlet-session invalidate-on-scheme-change="true"/>
```

This option helps make your system less vulnerable to sniffing of the session id or leakage of sensitive data from pages using HTTPS to other pages using HTTP.

13.9. CAPTCHA

Though strictly not part of the security API, Seam provides a built-in CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) algorithm to prevent automated processes from interacting with your application.

13.9.1. Configuring the CAPTCHA Servlet

To get up and running, it is necessary to configure the Seam Resource Servlet, which will provide the Captcha challenge images to your pages. This requires the following entry in `web.xml`:

```
<servlet>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <servlet-class>org.jboss.seam.servlet.SeamResourceServlet</servlet-class>
</servlet>

<servlet-mapping>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <url-pattern>/seam/resource/*</url-pattern>
</servlet-mapping>
```

13.9.2. Adding a CAPTCHA to a form

Adding a CAPTCHA challenge to a form is extremely easy. Here's an example:

```
<h:graphicImage value="/seam/resource/captcha"/>
<h:inputText id="verifyCaptcha" value="#{captcha.response}" required="true">
  <s:validate />
</h:inputText>
<h:message for="verifyCaptcha"/>
```

That's all there is to it. The `graphicImage` control displays the CAPTCHA challenge, and the `inputText` receives the user's response. The response is automatically validated against the CAPTCHA when the form is submitted.

13.9.3. Customising the CAPTCHA algorithm

You may customize the CAPTCHA algorithm by overriding the built-in component:

```
@Name("org.jboss.seam.captcha")
@Scope(SESSION)
public class HitchhikersCaptcha extends Captcha
{
  @Override @Create
  public void init()
  {
    setChallenge("What is the answer to life, the universe and everything?");
    setCorrectResponse("42");
  }

  @Override
  public BufferedImage renderChallenge()
  {
    BufferedImage img = super.renderChallenge();
```

```

        img.getGraphics().drawOval(5, 3, 60, 14); //add an obscuring decoration
        return img;
    }
}

```

13.10. Security Events

The following table describes a number of events (see Chapter 5, *Events, interceptors and exception handling*) raised by Seam Security.

Table 13.2. Security Events

Event Key	Description
<code>org.jboss.seam.security.loginSuccessful</code>	Raised when a login attempt is successful.
<code>org.jboss.seam.security.loginFailed</code>	Raised when a login attempt fails.
<code>org.jboss.seam.security.notLoggedIn</code>	Raised when a security check fails when the user is not logged in.
<code>org.jboss.seam.security.notAuthorized</code>	Raised when a security check fails when the user is logged in however doesn't have sufficient privileges.
<code>org.jboss.seam.security.preAuthenticate</code>	Raised just prior to user authentication.
<code>org.jboss.seam.security.postAuthenticate</code>	Raised just after user authentication.
<code>org.jboss.seam.security.loggedOut</code>	Raised after the user has logged out.
<code>org.jboss.seam.security.credentialsUpdated</code>	Raised when the user's credentials have been changed.
<code>org.jboss.seam.security.rememberMe</code>	Raised when the Identity's rememberMe property is changed.

13.11. Extending the Identity component

Sometimes it might be necessary to extend the Identity component if your application has special security requirements. For example, users might be required to authenticate using a Company or Department ID, along

with their usual username and password. If permission-based security is required then `RuleBasedIdentity` should be extended, otherwise `Identity` should be extended.

The following example shows an extended `Identity` component with an additional `companyCode` field. The install precedence of `APPLICATION` ensures that this extended `Identity` gets installed in preference to the built-in `Identity`.

```
@Name("org.jboss.seam.security.identity")
@Scope(SESSION)
@Install(precedence = APPLICATION)
@BypassInterceptors
@Startup
public class CustomIdentity extends Identity
{
    private static final LogProvider log = Logging.getLogProvider(CustomIdentity.class);

    private String companyCode;

    public String getCompanyCode()
    {
        return companyCode;
    }

    public void setCompanyCode(String companyCode)
    {
        this.companyCode = companyCode;
    }

    @Override
    public String login()
    {
        log.info("##### CUSTOM LOGIN CALLED #####");
        return super.login();
    }
}
```

Chapter 14. Internationalization and themes

Seam makes it easy to build internationalized applications by providing several built-in components for handling multi-language UI messages.

14.1. Locales

Each user login session has an associated instance of `java.util.Locale` (available to the application as a component named `locale`). Under normal circumstances, you won't need to do any special configuration to set the locale. Seam just delegates to JSF to determine the active locale:

- If there is a locale associated with the HTTP request (the browser locale), and that locale is in the list of supported locales from `faces-config.xml`, use that locale for the rest of the session.
- Otherwise, if a default locale was specified in the `faces-config.xml`, use that locale for the rest of the session.
- Otherwise, use the default locale of the server.

It is *possible* to set the locale manually via the Seam configuration properties `org.jboss.seam.international.localeSelector.language`, `org.jboss.seam.international.localeSelector.country` and `org.jboss.seam.international.localeSelector.variant`, but we can't think of any good reason to ever do this.

It is, however, useful to allow the user to set the locale manually via the application user interface. Seam provides built-in functionality for overriding the locale determined by the algorithm above. All you have to do is add the following fragment to a form in your JSP or Facelets page:

```
<h:selectOneMenu value="#{localeSelector.language}">
  <f:selectItem itemLabel="English" itemValue="en"/>
  <f:selectItem itemLabel="Deutsch" itemValue="de"/>
  <f:selectItem itemLabel="Francais" itemValue="fr"/>
</h:selectOneMenu>
<h:commandButton action="#{localeSelector.select}" value="#{messages['ChangeLanguage']}"/>
```

Or, if you want a list of all supported locales from `faces-config.xml`, just use:

```
<h:selectOneMenu value="#{localeSelector.localeString}">
  <f:selectItems value="#{localeSelector.supportedLocales}"/>
</h:selectOneMenu>
<h:commandButton action="#{localeSelector.select}" value="#{messages['ChangeLanguage']}"/>
```

When this user selects an item from the drop-down, and clicks the button, the Seam and JSF locales will be overridden for the rest of the session.

14.2. Labels

JSF supports internationalization of user interface labels and descriptive text via the use of `<f:loadBundle />`. You can use this approach in Seam applications. Alternatively, you can take advantage of the Seam messages component to display templated labels with embedded EL expressions.

14.2.1. Defining labels

Seam provides a `java.util.ResourceBundle` (available to the application as a `org.jboss.seam.core.resourceBundle`). You'll need to make your internationalized labels available via this special resource bundle. By default, the resource bundle used by Seam is named `messages` and so you'll need to define your labels in files named `messages.properties`, `messages_en.properties`, `messages_en_AU.properties`, etc. These files usually belong in the `WEB-INF/classes` directory.

So, in `messages_en.properties`:

```
Hello=Hello
```

And in `messages_en_AU.properties`:

```
Hello=G'day
```

You can select a different name for the resource bundle by setting the Seam configuration property named `org.jboss.seam.core.resourceLoader.bundleNames`. You can even specify a list of resource bundle names to be searched (depth first) for messages.

```
<core:resource-loader>
  <core:bundle-names>
    <value>mycompany_messages</value>
    <value>standard_messages</value>
  </core:bundle-names>
</core:resource-loader>
```

If you want to define a message just for a particular page, you can specify it in a resource bundle with the same name as the JSF view id, with the leading `/` and trailing file extension removed. So we could put our message in `welcome/hello_en.properties` if we only needed to display the message on `/welcome/hello.jsp`.

You can even specify an explicit bundle name in `pages.xml`:

```
<page view-id="/welcome/hello.jsp" bundle="HelloMessages"/>
```

Then we could use messages defined in `HelloMessages.properties` on `/welcome/hello.jsp`.

14.2.2. Displaying labels

If you define your labels using the Seam resource bundle, you'll be able to use them without having to type `<f:loadBundle ... />` on every page. Instead, you can simply type:

```
<h:outputText value="#{messages['Hello']}" />
```

or:

```
<h:outputText value="#{messages.Hello}" />
```

Even better, the messages themselves may contain EL expressions:

```
Hello=Hello, #{user.firstName} #{user.lastName}
```

```
Hello=G'day, #{user.firstName}
```

You can even use the messages in your code:

```
@In private Map<String, String> messages;
```

```
@In("#{messages['Hello']}") private String helloMessage;
```

14.2.3. Faces messages

The `facesMessages` component is a super-convenient way to display success or failure messages to the user. The functionality we just described also works for faces messages:

```
@Name("hello")
@Stateless
public class HelloBean implements Hello {
    @In FacesMessages facesMessages;

    public String sayIt() {
        facesMessages.addFromResourceBundle("Hello");
    }
}
```

This will display `Hello`, `Gavin King` or `G'day, Gavin`, depending upon the user's locale.

14.3. Timezones

There is also a session-scoped instance of `java.util.Timezone`, named `org.jboss.seam.international.timezone`, and a Seam component for changing the timezone named `org.jboss.seam.international.timezoneSelector`. By default, the timezone is the default timezone of the server. Unfortunately, the JSF specification says that all dates and times should be assumed to be UTC, and displayed as UTC, unless a timezone is explicitly specified using `<f:convertDateTime>`. This is an extremely inconvenient default behavior.

Seam overrides this behavior, and defaults all dates and times to the Seam timezone. In addition, Seam provides the `<s:convertDateTime>` tag which always performs conversions in the Seam timezone.

14.4. Themes

Seam applications are also very easily skinnable. The theme API is very similar to the localization API, but of course these two concerns are orthogonal, and some applications support both localization and themes.

First, configure the set of supported themes:

```
<theme:theme-selector cookie-enabled="true">
    <theme:available-themes>
        <value>default</value>
        <value>accessible</value>
        <value>printable</value>
    </theme:available-themes>
</theme:theme-selector>
```

Note that the first theme listed is the default theme.

Themes are defined in a properties file with the same name as the theme. For example, the default theme is defined as a set of entries in `default.properties`. For example, `default.properties` might define:

```
css ../screen.css
template /template.xhtml
```

Usually the entries in a theme resource bundle will be paths to CSS styles or images and names of facelets templates (unlike localization resource bundles which are usually text).

Now we can use these entries in our JSP or facelets pages. For example, to theme the stylesheet in a facelets page:

```
<link href="#{theme.css}" rel="stylesheet" type="text/css" />
```

Or, when the page definition resides in a subdirectory:

```
<link href="#{facesContext.externalContext.requestContextPath}#{theme.css}"
      rel="stylesheet" type="text/css" />
```

Most powerfully, facelets lets us theme the template used by a `<ui:composition>`:

```
<ui:composition xmlns="http://www.w3.org/1999/xhtml"
  xmlns:ui="http://java.sun.com/jsf/facelets"
  xmlns:h="http://java.sun.com/jsf/html"
  xmlns:f="http://java.sun.com/jsf/core"
  template="#{theme.template}">
```

Just like the locale selector, there is a built-in theme selector to allow the user to freely switch themes:

```
<h:selectOneMenu value="#{themeSelector.theme}">
  <f:selectItems value="#{themeSelector.themes}" />
</h:selectOneMenu>
<h:commandButton action="#{themeSelector.select}" value="Select Theme"/>
```

14.5. Persisting locale and theme preferences via cookies

The locale selector, theme selector and timezone selector all support persistence of locale and theme preference to a cookie. Simply set the `cookie-enabled` property in `components.xml`:

```
<theme:theme-selector cookie-enabled="true">
  <theme:available-themes>
    <value>default</value>
    <value>accessible</value>
    <value>printable</value>
  </theme:available-themes>
</theme:theme-selector>

<international:locale-selector cookie-enabled="true"/>
```

Chapter 15. Seam Text

Collaboration-oriented websites require a human-friendly markup language for easy entry of formatted text in forum posts, wiki pages, blogs, comments, etc. Seam provides the `<s:formattedText/>` control for display of formatted text that conforms to the *Seam Text* language. Seam Text is implemented using an ANTLR-based parser. You don't need to know anything about ANTLR to use it, however.

15.1. Basic fomattting

Here is a simple example:

```
It's easy to make *bold text*, /italic text/, |monospace|,
~deleted text~, super^scripts^ or _underlines_.
```

If we display this using `<s:formattedText/>`, we will get the following HTML produced:

```
<p>
It's easy to make <b>bold text</b>, <i>italic text</i>, <tt>monospace</tt>
<del>deleted text</del>, super<sup>scripts</sup> or <u>underlines</u>.
</p>
```

We can use a blank line to indicate a new paragraph, and `+` to indicate a heading:

```
+This is a big heading
You /must/ have some text following a heading!

++This is a smaller heading
This is the first paragraph. We can split it across multiple
lines, but we must end it with a blank line.

This is the second paragraph.
```

(Note that a simple newline is ignored, you need an additional blank line to wrap text into a new paragraph.) This is the HTML that results:

```
<h1>This is a big heading</h1>
<p>
You <i>must</i> have some text following a heading!
</p>

<h2>This is a smaller heading</h2>
<p>
This is the first paragraph. We can split it across multiple
lines, but we must end it with a blank line.
</p>

<p>
This is the second paragraph.
</p>
```

Ordered lists are created using the `#` character. Unordered lists use the `=` character:

```
An ordered list:

#first item
#second item
#and even the /third/ item

An unordered list:
```

```
=an item
=another item
```

```
<p>
An ordered list:
</p>

<ol>
<li>first item</li>
<li>second item</li>
<li>and even the <i>third</i> item</li>
</ol>

<p>
An unordered list:
</p>

<ul>
<li>an item</li>
<li>another item</li>
</ul>
```

Quoted sections should be surrounded in double quotes:

```
The other guy said:

"Nyeah nyeah-nee
/nyeah/ nyeah!"

But what do you think he means by "nyeah-nee"?
```

```
<p>
The other guy said:
</p>

<q>Nyeah nyeah-nee
<i>nyeah</i> nyeah!</q>

<p>
But what do you think he means by <q>nyeah-nee</q>?
</p>
```

15.2. Entering code and text with special characters

Special characters such as *, | and #, along with HTML characters such as <, > and & may be escaped using \:

```
You can write down equations like 2\*3\=6 and HTML tags
like \<body\> using the escape character: \.
```

```
<p>
You can write down equations like 2*3=6 and HTML tags
like &lt;body&gt; using the escape character: \.
</p>
```

And we can quote code blocks using backticks:

```
My code doesn't work:

`for (int i=0; i<100; i--)
{
    doSomething();
}`
```

```
Any ideas?
```

```
<p>
My code doesn't work:
</p>

<pre>for (int i=0; i<100; i--)
{
    doSomething();
}</pre>

<p>
Any ideas?
</p>
```

Note that inline monospace formatting always escapes (most monospace formatted text is in fact code or tags with many special characters). So you can, for example, write:

```
This is a |<tag attribute="value"/>| example.
```

without escaping any of the characters inside the monospace bars. The downside is that you can't format inline monospace text in any other way (italics, underscore, and so on).

15.3. Links

A link may be created using the following syntax:

```
Go to the Seam website at [=http://jboss.com/products/seam].
```

Or, if you want to specify the text of the link:

```
Go to [the Seam website=>http://jboss.com/products/seam].
```

For advanced users, it is even possible to customize the Seam Text parser to understand wikiword links written using this syntax.

15.4. Entering HTML

Text may even include a certain limited subset of HTML (don't worry, the subset is chosen to be safe from cross-site scripting attacks). This is useful for creating links:

```
You might want to link to <a href="http://jboss.com/products/seam">something
cool</a>, or even include an image: 
```

And for creating tables:

```
<table>
  <tr><td>First name:</td><td>Gavin</td></tr>
  <tr><td>Last name:</td><td>King</td></tr>
</table>
```

But you can do much more if you want!

Chapter 16. iText PDF generation

Seam now includes an component set for generating documents using iText. The primary focus of Seam's iText document support is for the generation of PDF documents, but Seam also offers basic support for RTF document generation.

16.1. Using PDF Support

iText support is provided by `jboss-seam-pdf.jar`. This JAR contains the iText JSF controls, which are used to construct views that can render to PDF, and the `DocumentStore` component, which serves the rendered documents to the user. To include PDF support in your application, include `jboss-seam-pdf.jar` in your `WEB-INF/lib` directory along with the iText JAR file. There is no further configuration needed to use Seam's iText support.

The Seam iText module requires the use of Facelets as the view technology. Future versions of the library may also support the use of JSP. Additionally, it requires the use of the `seam-ui` package.

The `examples/itext` project contains an example of the PDF support in action. It demonstrates proper deployment packaging, and it contains a number examples that demonstrate the key PDF generation features currently supported.

16.1.1. Creating a document

<code><p:document></code>	<i>Description</i> Documents are generated by facelets documents using tags in the <code>http://jboss.com/products/seam/pdf</code> namespace. Documents should always have the document tag at the root of the document. The document tag prepares Seam to generate a document into the <code>DocumentStore</code> and renders an HTML redirect to that stored content. <i>Attributes</i> <code>type</code> — The type of the document to be produced. Valid values are PDF, RTF and HTML modes. Seam defaults to PDF generation, and many of the features only work correctly when generating PDF documents. <code>pageSize</code> — The size of the page to be generated. The most commonly used values would be LETTER and A4. A full list of supported page sizes can be found in <code>com.lowagie.text.PageSize</code> class. Alternatively, <code>pageSize</code> can provide the width and height of the page directly. The value "612 792", for example, is equivalent to the LETTER page size. <code>orientation</code> — The orientation of the page. Valid values are portrait and landscape. In landscape mode, the height and width page size values are reversed. <code>margins</code> — The left, right, top and bottom margin values. <code>marginMirroring</code> — Indicates that margin settings should be reversed an al-
---------------------------------	---

	ternating pages. <i>Metadata Attributes</i> • title • subject • keywords • author • creator <i>Usage</i> <pre><p:document xmlns:p="http://jboss.com/products/seam/pdf"> The document goes here. </p:document></pre>
--	--

16.1.2. Basic Text Elements

Useful documents will need to contain more than just text; however, the standard UI components are geared towards HTML generation and are not useful for generating PDF content. Instead, Seam provides a special UI components for generating suitable PDF content. Tags like `<p:image>` and `<p:paragraph>` are the basic foundations of simple documents. Tags like `<p:font>` provide style information to all the content surrounding them.

<code><p:paragraph></code>	<i>Description</i> Most uses of text should be sectioned into paragraphs so that text fragments can be flowed, formatted and styled in logical groups. <i>Attributes</i> • firstLineIndent • extraParagraphSpace • leading • multipliedLeading • spacingBefore — The blank space to be inserted before the element. • spacingAfter — The blank space to be inserted after the element. • indentationLeft • indentationRight • keepTogether
----------------------------------	---

	<i>Usage</i> <pre><p:paragraph alignment="justify"> This is a simple document. It isn't very fancy. </p:paragraph></pre>
<code><p:text></code>	<i>Description</i> The text tag allows text fragments to be produced from application data using normal JSF converter mechanisms. It is very similar to the <code>outputText</code> tag used when rendering HTML documents. <i>Attributes</i> value — The value to be displayed. This will typically be a value binding expression. <i>Usage</i> <pre><p:paragraph> The item costs <p:text value="#{product.price}"> <f:convertNumber type="currency" currencySymbol="\$"/> </p:text> </p:paragraph></pre>
<code><p:font></code>	<i>Description</i> The font tag defines the default font to be used for all text inside of it. <i>Attributes</i> - name — The font name, for example: COURIER, HELVETICA, TIMES-ROMAN, SYMBOL or ZAPFDINGBATS. - size — The point size of the font. - style — The font styles. Any combination of : NORMAL, BOLD, ITALIC, OBLIQUE, UNDERLINE, LINE-THROUGH - encoding — The character set encoding. <i>Usage</i> <pre><p:font family="courier" style="bold" size="24"> <p:paragraph>My Title</p:paragraph> </p:font></pre>
<code><p:newPage></code>	<i>Description</i>

	<code>p:newPage</code> inserts a page break. <i>Usage</i> <pre><p:newPage /></pre>
--	--

<code><p:image></code>	<i>Description</i> <code>p:image</code> inserts an image into the document. Images can be loaded from the classpath or from the web application context using the <code>value</code> attribute. Resources can also be dynamically generated by application code. The <code>imageData</code> attribute can specify a value binding expression whose value is a <code>java.awt.Image</code> object. <i>Attributes</i> • <code>value</code> — A resource name or a method expression binding to an application-generated image. • <code>rotation</code> — The rotation of the image in degrees. • <code>height</code> — The height of the image. • <code>width</code> — The width of the image. • <code>alignment</code> — The alignment of the image. (see Section 16.1.7.2, “Alignment Values” for possible values) • <code>alt</code> — Alternative text representation for the image. • <code>indentationLeft</code> • <code>indentationRight</code> • <code>spacingBefore</code> — The blank space to be inserted before the element. • <code>spacingAfter</code> — The blank space to be inserted after the element. • <code>widthPercentage</code> • <code>initialRotation</code> • <code>dpi</code> • <code>scalePercent</code> — The scaling factor (as a percentage) to use for the image. This can be expressed as a single percentage value or as two percentage values representing separate x and y scaling percentages. • <code>wrap</code> • <code>underlying</code>
------------------------------	--

	<i>Usage</i> <pre><p:image value="/jboss.jpg" /></pre> <pre><p:image value="#{images.chart}" /></pre>
<p:anchor>	<i>Description</i> p:anchor defines clickable links from a document. It supports the following attributes: <i>Attributes</i> - name — The name of an in-document anchor destination. - reference — The destination the link refers to. Links to other points in the document should begin with a "#". For example, "#link1" to refer to an anchor position with a name of link1. Links may also be a full URL to point to a resource outside of the document. <i>Usage</i> <pre><p:listItem><p:anchor reference="#reason1">Reason 1</p:anchor></p:listItem> ... <p:paragraph> <p:anchor name="reason1">It's the quickest way to get "rich"</p:anchor> ... </p:paragraph></pre>

16.1.3. Headers and Footers

<p:header> <p:footer>	<i>Description</i> The p:header and p:footer components provide the ability to place header and footer text on each page of a generated document, with the exception of the first page. Header and footer declarations should appear near the top of a document. <i>Attributes</i> - alignment — The alignment of the header/footer box section. (see Section 16.1.7.2, “Alignment Values” for alignment values) - backgroundColor — The background color of the header/footer box. (see Section 16.1.7.1, “Color Values” for color values) - borderColor — The border color of the header/footer box. Individual border sides can be set using borderColorLeft, borderColorRight, borderColorTop and borderColorBottom.(see Section 16.1.7.1, “Color Values” for color values)
---	--

	• borderWidth — The width of the border. Individual border sides can be specified using borderWidthLeft, borderWidthRight, borderWidthTop and borderWidthBottom. <i>Usage</i> <pre><p:facet name="header"> <p:font size="12"> <p:footer borderWidthTop="1" borderColorTop="blue" borderWidthBottom="0" alignment="center"> Why Seam? [<p:pageNumber />] </p:footer> </p:font> </f:facet></pre>
--	--

<code><p:pageNumber></code>	<i>Description</i> The current page number can be placed inside of a header or footer using the <code>p:pageNumber</code> tag. The page number tag can only be used in the context of a header or footer and can only be used once. <i>Usage</i> <pre><p:footer borderWidthTop="1" borderColorTop="blue" borderWidthBottom="0" alignment="center"> Why Seam? [<p:pageNumber />] </p:footer></pre>
-----------------------------------	--

16.1.4. Chapters and Sections

<code><p:chapter></code> <code><p:section></code>	<i>Description</i> If the generated document follows a book/article structure, the <code>p:chapter</code> and <code>p:section</code> tags can be used to provide the necessary structure. Sections can only be used inside of chapters, but they may be nested arbitrarily deep. Most PDF viewers provide easy navigation between chapters and sections in a document. <i>Attributes</i> • alignment — The alignment of the header/footer box section. (see Section 16.1.7.2, “Alignment Values” for alignment values) • number — The chapter number. Every chapter should be assigned a chapter number. • numberDepth — The depth of numbering for section. All sections are numbered relative to their surrounding chapter/sections. The fourth section of the first section of chapter three would be section 3.1.4, if displayed at the default number depth of three. To omit the chapter number, a number depth of 2 should be used. In that case, the section number would be displayed as 1.4.
--	--

	<i>Usage</i> <pre><p:document xmlns:p="http://jboss.com/products/seam/pdf" title="Hello"> <p:chapter number="1"> <p:title><p:paragraph>Hello</p:paragraph></p:title> <p:paragraph>Hello #{user.name}!</p:paragraph> </p:chapter> <p:chapter number="2"> <p:title><p:paragraph>Goodbye</p:paragraph></p:title> <p:paragraph>Goodbye #{user.name}.</p:paragraph> </p:chapter> </p:document></pre>
<p:header>	<i>Description</i> Any chapter or section can contain a <code>p:title</code>. The title will be displayed next to the chapter/section number. The body of the title may contain raw text or may be a <code>p:paragraph</code>.

16.1.5. Lists

List structures can be displayed using the `p:list` and `p:listItem` tags. Lists may contain arbitrarily-nested sublists. List items may not be used outside of a list. The following document uses the `ui:repeat` tag to display a list of values retrieved from a Seam component.

```
<p:document xmlns:p="http://jboss.com/products/seam/pdf"
  xmlns:ui="http://java.sun.com/jsf/facelets"
  title="Hello">
  <p:list style="numbered">
    <ui:repeat value="#{documents}" var="doc">
      <p:listItem>#{doc.name}</p:listItem>
    </ui:repeat>
  </p:list>
</p:document>
```

<p:list>	<i>Attributes</i> <code>style</code> — The ordering/bulleting style of list. One of: <code>NUMBERED</code>, <code>LETTERED</code>, <code>GREEK</code>, <code>ROMAN</code>, <code>ZAPFDINGBATS</code>, <code>ZAPFDINGBATS_NUMBER</code>. If no style is given, the list items are bulleted. <code>listSymbol</code> — For bulleted lists, specifies the bullet symbol. <code>indent</code> — The indentation level of the list. <code>lowerCase</code> — For list styles using letters, indicates whether the letters should be lower case. <code>charNumber</code> — For <code>ZAPFDINGBATS</code>, indicates the character code of the
----------	--

	bullet character. • <code>numberType</code> — For <code>ZAPFDINGBATS_NUMBER</code>, indicates the numbering style. <i>Usage</i> <pre><p:list style="numbered"> <ui:repeat value="#{documents}" var="doc"> <p:listItem>#{doc.name}</p:listItem> </ui:repeat> </p:list></pre>
--	---

<code><p:listItem></code>	<i>Description</i> <code>p:listItem</code> supports the following attributes: <i>Attributes</i> • <code>alignment</code> — The alignment of the header/footer box section. (see Section 16.1.7.2, “Alignment Values” for alignment values) • <code>alignment</code> — The alignment of the list item. (See Section 16.1.7.2, “Alignment Values” for possible values) • <code>indentationLeft</code> — The left indentation amount. • <code>indentationRight</code> — The right indentation amount. • <code>listSymbol</code> — Overrides the default list symbol for this list item. <i>Usage</i> <pre>...</pre>
---------------------------------	--

16.1.6. Tables

Table structures can be created using the `p:table` and `p:cell` tags. Unlike many table structures, there is no explicit row declaration. If a table has 3 columns, then every 3 cells will automatically form a row. Header and footer rows can be declared, and the headers and footers will be repeated in the event a table structure spans multiple pages.

<code><p:table></code>	<i>Description</i> <code>p:table</code> supports the following attributes. <i>Attributes</i> • <code>columns</code> — The number of columns (cells) that make up a table row.
------------------------------	--

	• <code>widths</code> — The relative widths of each column. There should be one value for each column. For example: <code>widths="2 1 1"</code> would indicate that there are 3 columns and the first column should be twice the size of the second and third column. • <code>headerRows</code> — The initial number of rows which are considered to be headers or footer rows and should be repeated if the table spans multiple pages. • <code>footerRows</code> — The number of rows that are considered to be footer rows. This value is subtracted from the <code>headerRows</code> value. If document has 2 rows which make up the header and one row that makes up the footer, <code>headerRows</code> should be set to 3 and <code>footerRows</code> should be set to 1 • <code>widthPercentage</code> — The percentage of the page width that the table spans. • <code>horizontalAlignment</code> — The horizontal alignment of the table. (See Section 16.1.7.2, “Alignment Values” for possible values) • <code>skipFirstHeader</code> • <code>runDirection</code> • <code>lockedWidth</code> • <code>splitRows</code> • <code>spacingBefore</code> — The blank space to be inserted before the element. • <code>spacingAfter</code> — The blank space to be inserted after the element. • <code>extendLastRow</code> • <code>headersInEvent</code> • <code>splitLate</code> • <code>keepTogether</code> <i>Usage</i> <pre><p:table columns="3" headerRows="1"> <p:cell>name</p:cell> <p:cell>owner</p:cell> <p:cell>size</p:cell> <ui:repeat value="#{documents}" var="doc"> <p:cell>#{doc.name}</p:cell> <p:cell>#{doc.user.name}</p:cell> <p:cell>#{doc.size}</p:cell> </ui:repeat> </p:table></pre>
<code><p:cell></code>	<i>Description</i> <code>p:cell</code> supports the following attributes.

Attributes

- `colspan` — Cells can span more than one column by declaring a `colspan` greater than 1. Tables do not have the ability to span across multiple rows.
- `horizontalAlignment` — The horizontal alignment of the cell. (see Section 16.1.7.2, “Alignment Values” for possible values)
- `verticalAlignment` — The vertical alignment of the cell. (see Section 16.1.7.2, “Alignment Values” for possible values)
- `padding` — Padding on a given side can also be specified using `paddingLeft`, `paddingRight`, `paddingTop` and `paddingBottom`.
- `useBorderPadding`
- `leading`
- `multipliedLeading`
- `indent`
- `verticalAlignment`
- `extraParagraphSpace`
- `fixedHeight`
- `noWrap`
- `minimumHeight`
- `followingIndent`
- `rightIndent`
- `spaceCharRatio`
- `runDirection`
- `arabicOptions`
- `useAscender`
- `grayFill`
- `rotation`

Usage

```
<p:cell>...</p:cell>
```

16.1.7. Document Constants

This section documents some of the constants shared by attributes on multiple tags.

16.1.7.1. Color Values

Seam documents do not yet support a full color specification. Currently, only named colors are supported. They are: white, gray, lightgray, darkgray, black, red, pink, yellow, green, magenta, cyan and blue.

16.1.7.2. Alignment Values

Where alignment values are used, the Seam PDF supports the following horizontal alignment values: left, right, center, justify and justifyall. The vertical alignment values are top, middle, bottom, and baseline.

16.1.8. Configuring iText

Document generation works out of the box with no additional configuration needed. However, there are a few points of configuration that are needed for more serious applications.

The default implementation serves PDF documents from a generic URL, `/seam-doc.seam`. Many browsers (and users) would prefer to see URLs that contain the actual PDF name like `/myDocument.pdf`. This capability requires some configuration. To serve PDF files, all `*.pdf` resources should be mapped to the Seam Servlet Filter and to the DocumentStoreServlet:

```
<filter>
  <filter-name>Seam Servlet Filter</filter-name>
  <filter-class>org.jboss.seam.servlet.SeamServletFilter</filter-class>
</filter>

<filter-mapping>
  <filter-name>Seam Servlet Filter</filter-name>
  <url-pattern>*.pdf</url-pattern>
</filter-mapping>

<servlet>
  <servlet-name>Document Store Servlet</servlet-name>
  <servlet-class>org.jboss.seam.pdf.DocumentStoreServlet</servlet-class>
</servlet>

<servlet-mapping>
  <servlet-name>Document Store Servlet</servlet-name>
  <url-pattern>*.pdf</url-pattern>
</servlet-mapping>
```

The `useExtensions` option on the document store component completes the functionality by instructing the document store to generate URLs with the correct filename extension for the document type being generated.

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:pdf="http://jboss.com/products/seam/pdf">
  <pdf:documentStore useExtensions="true" />
</components>
```

Generated documents are stored in conversation scope and will expire when the conversation ends. At that point, references to the document will be invalid. To You can specify a default view to be shown when a document does not exist using the `errorPage` property of the `documentStore`.

```
<pdf:documentStore useExtensions="true" errorPage="/pdfMissing.seam" />
```

16.2. Charting

Charting support is also provided with `jboss-seam-pdf.jar`. Charts can be used in PDF documents or can be used as images in an HTML page. Charting requires the JFreeChart library (`jfreechart.jar` and `jcommon.jar`) to be added to the `WEB-INF/lib` directory. Three types of charts are currently supported: pie charts, bar charts and line charts.

<code><p:barchart></code>	<i>Description</i> Displays a bar chart. <i>Attributes</i> • <code>borderVisible</code> — Controls whether or not a border is displayed around the entire chart. • <code>borderPaint</code> — The color of the border, if visible; • <code>borderBackgroundPaint</code> — The default background color of the chart. • <code>borderStroke</code> — • <code>domainAxisLabel</code> — The text label for the domain axis. • <code>domainAxisPaint</code> — The color of the domain axis label. • <code>domainGridlinesVisible</code>— Controls whether or not gridlines for the domain axis are shown on the chart. • <code>domainGridlinePaint</code>— The color of the domain gridlines, if visible. • <code>domainGridlineStroke</code> — The stroke style of the domain gridleines, if visible. • <code>height</code> — The height of the chart. • <code>width</code> — The width of the chart. • <code>is3D</code> — A boolean value indicating that the chart should be rendered in 3D instead of 2D. • <code>legend</code> — A boolean value indicating whether or not the chart should include a legend. • <code>legendItemPaint</code>— The default color of the text labels in the legend. • <code>legendItemBackgroundPaint</code>— The background color for the legend, if different from the chart background color. • <code>orientation</code> — The orientation of the plot, either vertical (the default) or horizontal. • <code>plotBackgroundPaint</code>— The color of the plot background.
---------------------------------	---

	• <code>plotBackgroundAlpha</code>— The alpha (transparency) level of the plot background. It should be a number between 0 (completely transparent) and 1 (completely opaque). • <code>plotForegroundAlpha</code>— The alpha (transparency) level of the plot. It should be a number between 0 (completely transparent) and 1 (completely opaque). • <code>plotOutlinePaint</code>— The color of the range gridlines, if visible. • <code>plotOutlineStroke</code> — The stroke style of the range gridlines, if visible. • <code>rangeAxisLabel</code> — The text label for the range axis. • <code>rangeAxisPaint</code> — The color of the range axis label. • <code>rangeGridlinesVisible</code>— Controls whether or not gridlines for the range axis are shown on the chart. • <code>rangeGridlinePaint</code>— The color of the range gridlines, if visible. • <code>rangeGridlineStroke</code> — The stroke style of the range gridlines, if visible. • <code>title</code> — The chart title text. • <code>titlePaint</code>— The color of the chart title text. • <code>titleBackgroundPaint</code>— The background color around the chart title. • <code>width</code> — The width of the chart. <i>Usage</i> <pre><p:barchart title="Bar Chart" legend="true" width="500" height="500"> <p:series key="Last Year"> <p:data columnKey="Joe" value="100" /> <p:data columnKey="Bob" value="120" /> </p:series> <p:series key="This Year"> <p:data columnKey="Joe" value="125" /> <p:data columnKey="Bob" value="115" /> </p:series> </p:barchart></pre>
--	--

<code><p:linechart></code>	<i>Description</i> Displays a line chart. <i>Attributes</i> • <code>borderVisible</code> — Controls whether or not a border is displayed around the entire chart. • <code>borderPaint</code> — The color of the border, if visible;
----------------------------------	---

- `borderBackgroundPaint` — The default background color of the chart.
- `borderStroke` —
- `domainAxisLabel` — The text label for the domain axis.
- `domainAxisPaint` — The color of the domain axis label.
- `domainGridlinesVisible`— Controls whether or not gridlines for the domain axis are shown on the chart.
- `domainGridlinePaint`— The color of the domain gridlines, if visible.
- `domainGridlineStroke` — The stroke style of the domain gridlines, if visible.
- `height` — The height of the chart.
- `width` — The width of the chart.
- `is3D` — A boolean value indicating that the chart should be rendered in 3D instead of 2D.
- `legend` — A boolean value indicating whether or not the chart should include a legend.
- `legendItemPaint`— The default color of the text labels in the legend.
- `legendItemBackgroundPaint`— The background color for the legend, if different from the chart background color.
- `orientation` — The orientation of the plot, either vertical (the default) or horizontal.
- `plotBackgroundPaint`— The color of the plot background.
- `plotBackgroundAlpha`— The alpha (transparency) level of the plot background. It should be a number between 0 (completely transparent) and 1 (completely opaque).
- `plotForegroundAlpha`— The alpha (transparency) level of the plot. It should be a number between 0 (completely transparent) and 1 (completely opaque).
- `plotOutlinePaint`— The color of the range gridlines, if visible.
- `plotOutlineStroke` — The stroke style of the range gridlines, if visible.
- `rangeAxisLabel` — The text label for the range axis.
- `rangeAxisPaint` — The color of the range axis label.
- `rangeGridlinesVisible`— Controls whether or not gridlines for the range axis are shown on the chart.
- `rangeGridlinePaint`— The color of the range gridlines, if visible.

	• <code>rangeGridlineStroke</code> — The stroke style of the range gridlines, if visible. • <code>title</code> — The chart title text. • <code>titlePaint</code> — The color of the chart title text. • <code>titleBackgroundPaint</code> — The background color around the chart title. • <code>width</code> — The width of the chart. <i>Usage</i> <pre><p:linechart title="Line Chart" width="500" height="500"> <p:series key="Prices"> <p:data columnKey="2003" value="7.36" /> <p:data columnKey="2004" value="11.50" /> <p:data columnKey="2005" value="34.625" /> <p:data columnKey="2006" value="76.30" /> <p:data columnKey="2007" value="85.05" /> </p:series> </p:linechart></pre>
--	---

<code><p:piechart></code>	<i>Description</i> Displays a pie chart. <i>Attributes</i> • <code>title</code> — The chart title text. • <code>label</code> — The default label text for pie sections. • <code>legend</code> — A boolean value indicating whether or not the chart should include a legend. Default value is true • <code>is3D</code> — A boolean value indicating that the chart should be rendered in 3D instead of 2D. • <code>labelLinkMargin</code> — The link margin for labels. • <code>labelLinkPaint</code> — The paint used for the label linking lines. • <code>labelLinkStroke</code> — the stroke used for the label linking lines. • <code>labelLinksVisible</code> — A flag that controls whether or not the label links are drawn. • <code>labelOutlinePaint</code> — The paint used to draw the outline of the section labels. • <code>labelOutlineStroke</code> — The stroke used to draw the outline of the section labels. • <code>labelShadowPaint</code> — The paint used to draw the shadow for the section la-
---------------------------------	--

	bels. • <code>labelPaint</code>— The color used to draw the section labels • <code>labelGap</code>— The gap between the labels and the plot as a percentage of the plot width. • <code>labelBackgroundPaint</code>— The color used to draw the background of the section labels. If this is null, the background is not filled. • <code>startAngle</code>— The starting angle of the first section. • <code>circular</code>— A boolean value indicating that the chart should be drawn as a circle. If false, the chart is drawn as an ellipse. The default is true. • <code>direction</code>— The direction the pie section are drawn. One of: <code>clockwise</code> or <code>anticlockwise</code>. The default is <code>clockwise</code>. • <code>sectionOutlinePaint</code>— The outline paint for all sections. • <code>sectionOutlineStroke</code>— The outline stroke for all sections • <code>sectionOutlinesVisible</code>— Indicates whether an outline is drawn for each section in the plot. • <code>baseSectionOutlinePaint</code>— The base section outline paint. • <code>baseSectionPaint</code>— The base section paint. • <code>baseSectionOutlineStroke</code>— The base section outline stroke. <i>Usage</i> <pre><p:piechart title="Pie Chart" circular="false" direction="anticlockwise" startAngle="30" labelGap="0.1" labelLinkPaint="red"> <p:series key="Prices"> <p:data key="2003" columnKey="2003" value="7.36" /> <p:data key="2004" columnKey="2004" value="11.50" /> <p:data key="2005" columnKey="2005" value="34.625" /> <p:data key="2006" columnKey="2006" value="76.30" /> <p:data key="2007" columnKey="2007" value="85.05" /> </p:series> </p:piechart></pre>
--	---

<code><p:series></code>	<i>Description</i> Category data can be broken down into series. The series tag is used to categorize a set of data with a series and apply styling to the entire series. <i>Attributes</i>
-------------------------------	--

	• <code>key</code> — The series name. • <code>seriesPaint</code> — The color of each item in the series • <code>seriesOutlinePaint</code> — The outline color for each item in the series. • <code>seriesOutlineStroke</code> — The stroke used to draw each item in the series. • <code>seriesVisible</code> — A boolean indicating if the series should be displayed. • <code>seriesVisibleInLegend</code> — A boolean indicating if the series should be listed in the legend. <i>Usage</i> <pre><p:series key="data1"> <ui:repeat value="#{data.pieData1}" var="item"> <p:data columnKey="#{item.name}" value="#{item.value}" /> </ui:repeat> </p:series></pre>
--	---

<code><p:data></code>	<i>Description</i> The data tag describes each data point to be displayed in the graph. <i>Attributes</i> • <code>key</code> — The name of the data item. • <code>series</code> — The series name, when not embedded inside a <code><p:series></code>. • <code>value</code> — The numeric data value. • <code>explodedPercent</code> — For pie charts, indicates how exploded a from the pie a piece is. • <code>sectionOutlinePaint</code> — For bar charts, the color of the section outline. • <code>sectionOutlineStroke</code> — For bar charts, the stroke type for the section outline. • <code>sectionPaint</code> — For bar charts, the color of the section. <i>Usage</i> <pre><p:data key="foo" value="20" sectionPaint="#111111" explodedPercent=".2" /> <p:data key="bar" value="30" sectionPaint="#333333" /> <p:data key="baz" value="40" sectionPaint="#555555" sectionOutlineStroke="my-dot-style" /></pre>
-----------------------------	--

<code><p:color></code>	<i>Description</i>
------------------------------	--------------------

	The color component declares a color or gradient than can be referenced when drawing filled shapes. <i>Attributes</i> • <code>color</code> — The color value. For gradient colors, this the starting color. Section 16.1.7.1, “Color Values” • <code>color2</code> — For gradient colors, this is the color that ends the gradient. • <code>point</code> — The co-ordinates where the gradient color begins. • <code>point2</code> — The co-ordinates where the gradient color ends. <i>Usage</i> <pre><p:color id="foo" color="#0ff00f"/> <p:color id="bar" color="#ff00ff" color2="#00ff00" point="50 50" point2="300 300"/></pre>
--	--

<code><p:stroke></code>	<i>Description</i> Describes a stroke used to draw lines in a chart. <i>Attributes</i> • <code>width</code> — The width of the stroke. • <code>cap</code> — The line cap type. Valid values are <code>butt</code>, <code>round</code> and <code>square</code> • <code>join</code> — The line join type. Valid values are <code>miter</code>, <code>round</code> and <code>bevel</code> • <code>miterLimit</code> — For miter joins, this value is the limit of the size of the join. • <code>dash</code> — The dash value sets the dash pattern to be used to draw the line. The space separated integers indicate the length of each alternating drawn and un-drawn segments. • <code>dashPhase</code> — The dash phase indicates the offset into the dash pattern that the the line should be drawn with. <i>Usage</i> <pre><p:stroke id="dot2" width="2" cap="round" join="bevel" dash="2 3" /></pre>
-------------------------------	---

16.3. Bar codes

...

<p:barCode>	<i>Description</i> <i>Attributes</i> • <code>type</code> — A barcode type supported by iText. Valid values include: EAN13, EAN8, UPCA, UPCE, SUPP2, SUPP5, POSTNET, PLANET, CODE128, CODE128_UCC, CODE128_RAW and CODABAR. • <code>code</code>— The value to be encoded by the barcode. • <code>xpos</code>— For PDFs, the absolute x position of the barcode on the page. • <code>ypos</code>— For PDFs, the absolute y position of the barcode on the page. • <code>rotDegrees</code> — For PDFs, the rotation factor of the barcode in degrees. • <code>barHeight</code> — The height of the bars in the barCode • <code>minBarWidth</code> — The minimum bar width. • <code>barMultiplier</code> — The bar multiplier for wide bars or the distance between bars for POSTNET and PLANET code. • <code>barColor</code> — The color to draw the bars. • <code>textColor</code> — The color of any text on the barcode. • <code>textSize</code> — The size of the barcode text, if any. <i>Usage</i> <pre><p:barCode type="code128" barHeight="80" textSize="20" code="(10)45566(17)040301" codeType="code128_ucc" altText="My BarCode" /></pre>
-------------	--

16.4. Further documentation

For further information on iText, see:

- iText Home Page [<http://www.lowagie.com/iText/>]
- iText in Action [<http://www.manning.com/lowagie/>]

Chapter 17. Email

Seam now includes an optional components for templating and sending emails.

Email support is provided by `jboss-seam-mail.jar`. This JAR contains the mail JSF controls, which are used to construct emails, and the `mailSession` manager component.

The `examples/mail` project contains an example of the email support in action. It demonstrates proper packaging, and it contains a number of example that demonstrate the key features currently supported.

You can also test your mail's using Seam's integration testing environment. See Section 31.3.2, “Integration Testing Seam Mail”.

17.1. Creating a message

You don't need to learn a whole new templating language to use Seam Mail — an email is just facelet!

```
<m:message xmlns="http://www.w3.org/1999/xhtml"
  xmlns:m="http://jboss.com/products/seam/mail"
  xmlns:h="http://java.sun.com/jsf/html">

  <m:from name="Peter" address="peter@example.com" />
  <m:to name="#{person.firstname} #{person.lastname}">#{person.address}</m:to>
  <m:subject>Try out Seam!</m:subject>

  <m:body>
    <p><h:outputText value="Dear #{person.firstname}" />,</p>
    <p>You can try out Seam by visiting
    <a href="http://labs.jboss.com/jbossseam">http://labs.jboss.com/jbossseam</a>.</p>
    <p>Regards,</p>
    <p>Pete</p>
  </m:body>

</m:message>
```

The `<m:message>` tag wraps the whole message, and tells Seam to start rendering an email. Inside the `<m:message>` tag we use an `<m:from>` tag to set who the message is from, a `<m:to>` tag to specify a sender (notice how we use EL as we would in a normal facelet), and a `<m:subject>` tag.

The `<m:body>` tag wraps the body of the email. You can use regular HTML tags inside the body as well as JSF components.

So, now you have your email template, how do you go about sending it? Well, at the end of rendering the `m:message` the `mailSession` is called to send the email, so all you have to do is ask Seam to render the view:

```
@In(create=true)
private Renderer renderer;

public void send() {
    try {
        renderer.render("/simple.xhtml");
        facesMessages.add("Email sent successfully");
    }
    catch (Exception e) {
        facesMessages.add("Email sending failed: " + e.getMessage());
    }
}
```

If, for example, you entered an invalid email address, then an exception would be thrown, which is caught and

then displayed to the user.

17.1.1. Attachments

Seam makes it easy to attach files to an email. It supports most of the standard java types used when working with files.

If you wanted to email the `jboss-seam-mail.jar`:

```
<m:attachment value="/WEB-INF/lib/jboss-seam-mail.jar"/>
```

Seam will load the file from the classpath, and attach it to the email. By default it would be attached as `jboss-seam-mail.jar`; if you wanted it to have another name you would just add the `fileName` attribute:

```
<m:attachment value="/WEB-INF/lib/jboss-seam-mail.jar" fileName="this-is-so-cool.jar"/>
```

You could also attach a `java.io.File`, a `java.net.URL`:

```
<m:attachment value="#{numbers}"/>
```

Or a `byte[]` or a `java.io.InputStream`:

```
<m:attachment value="#{person.photo}" contentType="image/png"/>
```

You'll notice that for a `byte[]` and a `java.io.InputStream` you need to specify the MIME type of the attachment (as that information is not carried as part of the file).

And it gets even better, you can attach a Seam generated PDF, or any standard JSF view, just by wrapping a `<m:attachment>` around the normal tags you would use:

```
<m:attachment fileName="tiny.pdf">
  <p:document>
    A very tiny PDF
  </p:document>
</m:attachment>
```

If you had a set of files you wanted to attach (for example a set of pictures loaded from a database) you can just use a `<ui:repeat>`:

```
<ui:repeat value="#{people}" var="person">
  <m:attachment value="#{person.photo}" contentType="image/jpeg" fileName="#{person.firstname}_{person.lastname}.jpg"/>
</ui:repeat>
```

And if you want to display an attached image inline:

```
<m:attachment
  value="#{person.photo}"
  contentType="image/jpeg"
  fileName="#{person.firstname}_{person.lastname}.jpg"
  status="personPhoto"
  disposition="inline" />

```

You may be wondering what `cid:#{...}` does. Well, the IETF specified that by putting this as the `src` for your image, the attachments will be looked at when trying to locate the image (the Content-ID's must match) — magic!

You must declare the attachment before trying to access the status object.

17.1.2. HTML/Text alternative part

Whilst most mail readers nowadays support HTML, some don't, so you can add a plain text alternative to your email body:

```
<m:body>
  <f:facet name="alternative">Sorry, your email reader can't show our fancy email,
  please go to http://labs.jboss.com/jbossseam to explore Seam.</f:facet>
</m:body>
```

17.1.3. Multiple recipients

Often you'll want to send an email to a group of recipients (for example your users). All of the recipient mail tags can be placed inside a `<ui:repeat>`:

```
<ui:repeat value="#{allUsers}" var="user">
  <m:to name="#{user.firstname} #{user.lastname}" address="#{user.emailAddress}" />
</ui:repeat>
```

17.1.4. Multiple messages

Sometimes, however, you need to send a slightly different message to each recipient (e.g. a password reset). The best way to do this is to place the whole message inside a `<ui:repeat>`:

```
<ui:repeat value="#{people}" var="p">
  <m:message>
    <m:from name="#{person.firstname} #{person.lastname}" address="#{person.address}" />
    <m:to name="#{p.firstname}" address="#{p.address}" />
    ...
  </m:message>
</ui:repeat>
```

17.1.5. Templating

The mail templating example shows that facelets templating Just Works with the Seam mail tags.

Our `template.xhtml` contains:

```
<m:message>
  <m:from name="Seam" address="do-not-reply@jboss.com" />
  <m:to name="#{person.firstname} #{person.lastname}" address="#{person.address}" />
  <m:subject>#{subject}</m:subject>
  <m:body>
    <html>
      <body>
        <ui:insert name="body">This is the default body, specified by the template.</ui:insert>
      </body>
    </html>
  </m:body>
</m:message>
```

Our `templating.xhtml` contains:

```
<ui:param name="subject" value="Templating with Seam Mail" />
<ui:define name="body">
```

```
<p>This example demonstrates that you can easily use <i>facelets templating</i> in email!</p>
</ui:define>
```

You can also use facelets source tags in your email, but you must place them in a jar in `WEB-INF/lib` - referencing the `.taglib.xml` from `web.xml` isn't reliable when using Seam Mail (if you send your mail asynchronously Seam Mail doesn't have access to the full JSF or Servlet context, and so doesn't know about `web.xml` configuration parameters).

If you do need more configure Facelets or JSF when sending mail, you'll need to override the `Renderer` component and do the configuration programmatically - only for advanced users!

17.1.6. Internationalisation

Seam supports sending internationalised messages. By default, the encoding provided by JSF is used, but this can be overridden on the template:

```
<m:message charset="UTF-8">
    ...
</m:message>
```

The body, subject and recipient (and from) name will be encoded. You'll need to make sure facelets uses the correct charset for parsing your pages by setting encoding of the template:

```
<?xml version="1.0" encoding="UTF-8"?>
```

17.1.7. Other Headers

Sometimes you'll want to add other headers to your email. Seam provides support for some (see Section 17.5, “Tags”). For example, we can set the importance of the email, and ask for a read receipt:

```
<m:message xmlns:m="http://jboss.com/products/seam/mail"
    importance="low"
    requestReadReceipt="true"/>
```

Otherwise you can add any header to the message using the `<m:header>` tag:

```
<m:header name="X-Sent-From" value="JBoss Seam"/>
```

17.2. Receiving emails

If you are using EJB then you can use a MDB (Message Driven Bean) to receive email. JBoss provides a JCA adaptor — `mail-ra.rar` — but the version distributed with JBoss AS has a number of limitations (and isn't bundled in some versions) therefore we recommend using the `mail-ra.rar` distributed with Seam is recommended (it's in the `mail` directory in the Seam bundle). `mail-ra.rar` should be placed in `$JBOSS_HOME/server/default/deploy`; if the version of JBoss AS you use already has this file, replace it.

You can configure it like this:

```
@MessageDriven(activationConfig={
    @ActivationConfigProperty(propertyName="mailServer", propertyValue="localhost"),
    @ActivationConfigProperty(propertyName="mailFolder", propertyValue="INBOX"),
    @ActivationConfigProperty(propertyName="storeProtocol", propertyValue="pop3"),
    @ActivationConfigProperty(propertyName="userName", propertyValue="seam"),
    @ActivationConfigProperty(propertyName="password", propertyValue="seam")
})
```

```

})
@ResourceAdapter("mail-ra.rar")
@Name("maillistener")
public class MailListenerMDB implements MailListener {

    @In(create=true)
    private OrderProcessor orderProcessor;

    public void onMessage(Message message) {
        // Process the message
        orderProcessor.process(message.getSubject());
    }
}

```

Each message received will cause `onMessage(Message message)` to be called. Most Seam annotations will work inside a MDB but you mustn't access the persistence context.

You can find more information on `mail-ra.rar` at <http://wiki.jboss.org/wiki/Wiki.jsp?page=InboundJavaMail>.

If you aren't using JBoss AS you can still use `mail-ra.rar` or you may find your application server includes a similar adapter.

17.3. Configuration

To include Email support in your application, include `jboss-seam-mail.jar` in your `WEB-INF/lib` directory. If you are using JBoss AS there is no further configuration needed to use Seam's email support. Otherwise you need to make sure you have the JavaMail API, an implementation of the JavaMail API present (the API and impl used in JBoss AS are distributed with seam as `lib/mail.jar`), and a copy of the Java Activation Framework (distributed with Seam as `lib/activation.jar`).

The Seam Email module requires the use of Facelets as the view technology. Future versions of the library may also support the use of JSP. Additionally, it requires the use of the `seam-ui` package.

The `mailSession` component uses JavaMail to talk to a 'real' SMTP server.

17.3.1. mailSession

A JavaMail Session may be available via a JNDI lookup if you are working in an JEE environment or you can use a Seam configured Session.

The `mailSession` component's properties are described in more detail in Section 28.8, "Mail-related components".

17.3.1.1. JNDI lookup in JBoss AS

The JBossAS `deploy/mail-service.xml` configures a JavaMail session binding into JNDI. The default service configuration will need altering for your network. <http://wiki.jboss.org/wiki/Wiki.jsp?page=JavaMail> describes the service in more detail.

```

<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:mail="http://jboss.com/products/seam/mail">

  <mail:mail-session session-jndi-name="java:/Mail"/>

</components>

```

Here we tell Seam to get the mail session bound to `java:/Mail` from JNDI.

17.3.1.2. Seam configured Session

A mail session can be configured via `components.xml`. Here we tell Seam to use `smtp.example.com` as the smtp server:

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:mail="http://jboss.com/products/seam/mail">

  <mail:mail-session host="smtp.example.com"/>

</components>
```

17.4. Meldware

Seam's mail examples use Meldware (from buni.org [<http://buni.org>]) as a mail server. Meldware is a groupware package that provides SMTP, POP3, IMAP, webmail, a shared calendar and an graphical admin tool; it's written as a JEE application so can be deployed onto JBoss AS alongside your Seam application.

The version of Meldware distributed with Seam (in the `mail/buni-meldware` folder) is specially tailored for development - mailboxes, users and aliases (email addresses) are created every time the application deploys. If you want to use Meldware in production you should install the latest release from buni.org [<http://buni.org>].

To create mailboxes, users and aliases, you can use the `meldware` component:

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:core="http://jboss.com/products/seam/core"
  xmlns:mail="http://jboss.com/products/seam/mail">

  <mail:mail-session host="smtp.example.com"/>

  <mail:meldware>
    <mail:users>
      <value>#{duke}</value>
      <value>#{root}</value>
    </mail:users>
  </mail:meldware>

  <mail:meldware-user name="duke" username="duke" password="duke">
    <mail:aliases>
      <value>duke@jboss.org</value>
      <value>duke@jboss.com</value>
    </mail:aliases>
  </mail:meldware-user name="root" username="root" password="root" administrator="true" />

</components>
```

Here we've created two users, duke, who has two email addresses and an administrator with the username root.

17.5. Tags

Emails are generated using tags in the `http://jboss.com/products/seam/mail` namespace. Documents should always have the message tag at the root of the message. The message tag prepares Seam to generate an email.

The standard templating tags of facelets can be used as normal. Inside the body you can use any JSF tag; if it requires access to external resources (stylesheets, javascript) then be sure to set the `urlBase`.

<m:message>

Root tag of a mail message

- `importance` — low, normal or high. By default normal, this sets the importance of the mail message.
- `precedence` — sets the precedence of the message (e.g. bulk).
- `requestReadReceipt` — by default false, if set, a read receipt request will be added, with the read receipt being sent to the `From:` address.
- `urlBase` — If set, the value is prepended to the `requestContextPath` allowing you to use components such as `<h:graphicImage>` in your emails.

<m:from>

Set's the From: address for the email. You can only have one of these per email.

- `name` — the name the email should come from.
- `address` — the email address the email should come from.

<m:replyTo>

Set's the Reply-to: address for the email. You can only have one of these per email.

- `address` — the email address the email should come from.

<m:to>

Add a recipient to the email. Use multiple `<m:to>` tags for multiple recipients. This tag can be safely placed inside a repeat tag such as `<ui:repeat>`.

- `name` — the name of the recipient.
- `address` — the email address of the recipient.

<m:cc>

Add a cc recipient to the email. Use multiple `<m:cc>` tags for multiple ccs. This tag can be safely placed inside a iterator tag such as `<ui:repeat>`.

- `name` — the name of the recipient.
- `address` — the email address of the recipient.

<m:bcc>

Add a bcc recipient to the email. Use multiple `<m:bcc>` tags for multiple bccs. This tag can be safely placed inside a repeat tag such as `<ui:repeat>`.

- `name` — the name of the recipient.
- `address` — the email address of the recipient.

<m:header>

Add a header to the email (e.g. `X-Sent-From: JBoss Seam`)

- `name` — The name of the header to add (e.g. `X-Sent-From`).
- `value` — The value of the header to add (e.g. `JBoss Seam`).

`<m:attachment>`

Add an attachment to the email.

- `value` — The file to attach:
 - `String` — A `String` is interpreted as a path to file within the classpath
 - `java.io.File` — An EL expression can reference a `File` object
 - `java.net.URL` — An EL expression can reference a `URL` object
 - `java.io.InputStream` — An EL expression can reference an `InputStream`. In this case both a `fileName` and a `contentType` must be specified.
 - `byte[]` — An EL expression can reference an `byte[]`. In this case both a `fileName` and a `contentType` must be specified.

If the `value` attribute is omitted:

- If this tag contains a `<p:document>` tag, the document described will be generated and attached to the email. A `fileName` should be specified.
- If this tag contains other JSF tags a HTML document will be generated from them and attached to the email. A `fileName` should be specified.
- `fileName` — Specify the file name to use for the attached file.
- `contentType` — Specify the MIME type of the attached file

`<m:subject>`

Set's the subject for the email.

`<m:body>`

Set's the body for the email. Supports an `alternative` facet which, if an HTML email is generated can contain alternative text for a mail reader which doesn't support html.

- `type` — If set to `plain` then a plain text email will be generated otherwise an HTML email is generated.

Chapter 18. Asynchronicity and messaging

Seam makes it very easy to perform work asynchronously from a web request. When most people think of asynchronicity in Java EE, they think of using JMS. This is certainly one way to approach the problem in Seam, and is the right way when you have strict and well-defined quality of service requirements. Seam makes it easy to send and receive JMS messages using Seam components.

But for many usecases, JMS is overkill. Seam layers a simple asynchronous method and event facility over your choice of *dispatchers*:

- `java.util.concurrent.ScheduledThreadPoolExecutor` (by default)
- the EJB timer service (for EJB 3.0 environments)
- Quartz

18.1. Asynchronicity

Asynchronous events and method calls have the same quality of service expectations as the underlying dispatcher mechanism. The default dispatcher, based upon a `ScheduledThreadPoolExecutor` performs efficiently but provides no support for persistent asynchronous tasks, and hence no guarantee that a task will ever actually be executed. If you're working in an environment that supports EJB 3.0, and add the following line to `components.xml`:

```
<async:timer-service-dispatcher/>
```

then your asynchronous tasks will be processed by the container's EJB timer service. If you're not familiar with the Timer service, don't worry, you don't need to interact with it directly if you want to use asynchronous methods in Seam. The important thing to know is that any good EJB 3.0 implementation will have the option of using persistent timers, which gives some guarantee that the tasks will eventually be processed.

Another alternative is to use the open source Quartz library to manage asynchronous method. You need to bundle the Quartz library JAR (found in the `lib` directory) in your EAR and declare it as a Java module in `application.xml`. In addition, you need to add the following line to `components.xml` to install the Quartz dispatcher.

```
<async:quartz-dispatcher/>
```

The Seam API for the default `ScheduledThreadPoolExecutor`, the EJB3 Timer, and the Quartz Scheduler are largely the same. They can just "plug and play" by adding a line to `components.xml`.

18.1.1. Asynchronous methods

In simplest form, an asynchronous call just lets a method call be processed asynchronously (in a different thread) from the caller. We usually use an asynchronous call when we want to return an immediate response to the client, and let some expensive work be processed in the background. This pattern works very well in applications which use AJAX, where the client can automatically poll the server for the result of the work.

For EJB components, we annotate the local interface to specify that a method is processed asynchronously.

```
@Local
```

```
public interface PaymentHandler
{
    @Asynchronous
    public void processPayment(Payment payment);
}
```

(For JavaBean components we can annotate the component implementation class if we like.)

The use of asynchronicity is transparent to the bean class:

```
@Stateless
@Name("paymentHandler")
public class PaymentHandlerBean implements PaymentHandler
{
    public void processPayment(Payment payment)
    {
        //do some work!
    }
}
```

And also transparent to the client:

```
@Stateful
@Name("paymentAction")
public class CreatePaymentAction
{
    @In(create=true) PaymentHandler paymentHandler;
    @In Bill bill;

    public String pay()
    {
        paymentHandler.processPayment( new Payment(bill) );
        return "success";
    }
}
```

The asynchronous method is processed in a completely new event context and does not have access to the session or conversation context state of the caller. However, the business process context is propagated.

Asynchronous method calls may be scheduled for later execution using the `@Duration`, `@Expiration` and `@IntervalDuration` annotations.

```
@Local
public interface PaymentHandler
{
    @Asynchronous
    public void processScheduledPayment(Payment payment, @Expiration Date date);

    @Asynchronous
    public void processRecurringPayment(Payment payment,
                                       @Expiration Date date,
                                       @IntervalDuration Long interval)'
}
```

```
@Stateful
@Name("paymentAction")
public class CreatePaymentAction
{
    @In(create=true) PaymentHandler paymentHandler;
    @In Bill bill;

    public String schedulePayment()
    {
        paymentHandler.processScheduledPayment( new Payment(bill), bill.getDueDate() );
        return "success";
    }
}
```

```

    }

    public String scheduleRecurringPayment()
    {
        paymentHandler.processRecurringPayment( new Payment(bill), bill.getDueDate(),
                                                ONE_MONTH );
        return "success";
    }
}

```

Both client and server may access the `Timer` object associated with the invocation. The `Timer` object shown below is the EJB3 timer when you use the EJB3 dispatcher. For the default `ScheduledThreadPoolExecutor`, the returned object is `Future` from the JDK. For the Quartz dispatcher, it returns `QuartzTriggerHandle`, which we will discuss in the next section.

```

@Local
public interface PaymentHandler
{
    @Asynchronous
    public Timer processScheduledPayment(Payment payment, @Expiration Date date);
}

```

```

@Stateless
@Name("paymentHandler")
public class PaymentHandlerBean implements PaymentHandler
{
    @In Timer timer;

    public Timer processScheduledPayment(Payment payment, @Expiration Date date)
    {
        //do some work!

        return timer; //note that return value is completely ignored
    }
}

```

```

@Stateful
@Name("paymentAction")
public class CreatePaymentAction
{
    @In(create=true) PaymentHandler paymentHandler;
    @In Bill bill;

    public String schedulePayment()
    {
        Timer timer = paymentHandler.processScheduledPayment( new Payment(bill),
                                                                bill.getDueDate() );
        return "success";
    }
}

```

Asynchronous methods cannot return any other value to the caller.

18.1.2. Asynchronous methods with the Quartz Dispatcher

The Quartz dispatcher (see earlier on how to install it) allows you to use the `@Asynchronous`, `@Duration`, `@Expiration`, and `@IntervalDuration` annotations as above. But it has some powerful additional features. The Quartz dispatcher supports three new annotations.

The `@FinalExpiration` annotation specifies an end date for the recurring task.

```
// Defines the method in the "processor" component
@Asynchronous
public QuartzTriggerHandle schedulePayment(@Expiration Date when,
                                           @IntervalDuration Long interval,
                                           @FinalExpiration Date endDate,
                                           Payment payment)
{
    // do the repeating or long running task until endDate
}

... ..

// Schedule the task in the business logic processing code
// Starts now, repeats every hour, and ends on May 10th, 2010
Calendar cal = Calendar.getInstance ();
cal.set (2010, Calendar.MAY, 10);
processor.schedulePayment(new Date(), 60*60*1000, cal.getTime(), payment);
```

Note that the method returns the `QuartzTriggerHandle` object, which you can use later to stop, pause, and resume the scheduler. The `QuartzTriggerHandle` object is serializable, so you can save it into the database if you need to keep it around for extended period of time.

```
QuartzTriggerHandle handle =
    processor.schedulePayment(payment.getPaymentDate(),
                             payment.getPaymentCron(),
                             payment);
payment.setQuartzTriggerHandle( handle );
// Save payment to DB

// later ...

// Retrieve payment from DB
// Cancel the remaining scheduled tasks
payment.getQuartzTriggerHandle().cancel();
```

The `@IntervalCron` annotation supports Unix cron job syntax for task scheduling. For instance, the following asynchronous method runs at 2:10pm and at 2:44pm every Wednesday in the month of March.

```
// Define the method
@Asynchronous
public QuartzTriggerHandle schedulePayment(@Expiration Date when,
                                           @IntervalCron String cron,
                                           Payment payment)
{
    // do the repeating or long running task
}

... ..

// Schedule the task in the business logic processing code
QuartzTriggerHandle handle =
    processor.schedulePayment(new Date(), "0 10,44 14 ? 3 WED", payment);
```

The `@IntervalBusinessDay` annotation supports invocation on the "nth Business Day" scenario. For instance, the following asynchronous method runs at 14:00 on the 2nd business day of each month. By default, it excludes all weekends and US federal holidays until 2010 from the business days.

```
// Define the method
@Asynchronous
public QuartzTriggerHandle schedulePayment(@Expiration Date when,
                                           @IntervalBusinessDay NthBusinessDay nth,
                                           Payment payment)
{
    // do the repeating or long running task
```

```

}

... ..

// Schedule the task in the business logic processing code
QuartzTriggerHandle handle =
    processor.schedulePayment(new Date(),
        new NthBusinessDay(2, "14:00", WEEKLY), payment);

```

The `NthBusinessDay` object contains the configuration of the invocation trigger. You can specify more holidays (e.g., company holidays, non-US holidays etc.) via the `additionalHolidays` property.

```

public class NthBusinessDay implements Serializable
{
    int n;
    String fireAtTime;
    List <Date> additionalHolidays;
    BusinessDayIntervalType interval;
    boolean excludeWeekends;
    boolean excludeUsFederalHolidays;

    public enum BusinessDayIntervalType { WEEKLY, MONTHLY, YEARLY }

    public NthBusinessDay ()
    {
        n = 1;
        fireAtTime = "12:00";
        additionalHolidays = new ArrayList <Date> ();
        interval = BusinessDayIntervalType.WEEKLY;
        excludeWeekends = true;
        excludeUsFederalHolidays = true;
    }
    ... ..
}

```

The `@IntervalDuration`, `@IntervalCron`, and `@IntervalNthBusinessDay` annotations are mutually exclusive. If they are used in the same method, a `RuntimeException` will be thrown.

18.1.3. Asynchronous events

Component-driven events may also be asynchronous. To raise an event for asynchronous processing, simply call the `raiseAsynchronousEvent()` method of the `Events` class. To schedule a timed event, call the `raiseTimedEvent()` method, passing a `schedule` object (for the default dispatcher or timer service dispatcher, use `TimerSchedule`). Components may observe asynchronous events in the usual way, but remember that only the business process context is propagated to the asynchronous thread.

18.2. Messaging in Seam

Seam makes it easy to send and receive JMS messages to and from Seam components.

18.2.1. Configuration

To configure Seam's infrastructure for sending JMS messages, you need to tell Seam about any topics and queues you want to send messages to, and also tell Seam where to find the `QueueConnectionFactory` and/or `TopicConnectionFactory`.

Seam defaults to using `UIL2ConnectionFactory` which is the usual connection factory for use with JBossMQ.

If you are using some other JMS provider, you need to set one or both of `queueConnectionFactoryJndiName` and `topicConnectionFactoryJndiName` in `seam.properties`, `web.xml` or `components.xml`.

You also need to list topics and queues in `components.xml` to install Seam managed `TopicPublishers` and `QueueSenders`:

```
<jms:managed-topic-publisher name="stockTickerPublisher"
                             auto-create="true"
                             topic-jndi-name="topic/stockTickerTopic"/>

<jms:managed-queue-sender name="paymentQueueSender"
                          auto-create="true"
                          queue-jndi-name="queue/paymentQueue"/>
```

18.2.2. Sending messages

Now, you can inject a JMS `TopicPublisher` and `TopicSession` into any component:

```
@In
private TopicPublisher stockTickerPublisher;
@In
private TopicSession topicSession;

public void publish(StockPrice price) {
    try
    {
        stockTickerPublisher.publish( topicSession.createObjectMessage(price) );
    }
    catch (Exception ex)
    {
        throw new RuntimeException(ex);
    }
}
```

Or, for working with a queue:

```
@In
private QueueSender paymentQueueSender;
@In
private QueueSession queueSession;

public void publish(Payment payment) {
    try
    {
        paymentQueueSender.send( queueSession.createObjectMessage(payment) );
    }
    catch (Exception ex)
    {
        throw new RuntimeException(ex);
    }
}
```

18.2.3. Receiving messages using a message-driven bean

You can process messages using any EJB3 message driven bean. Message-driven beans may even be Seam components, in which case it is possible to inject other event and application scoped Seam components.

18.2.4. Receiving messages in the client

Seam Remoting lets you subscribe to a JMS topic from client-side JavaScript. This is described in Chapter 21, *Remoting*.

Chapter 19. Caching

In almost all enterprise applications, the database is the primary bottleneck, and the least scalable tier of the runtime environment. People from a PHP/Ruby environment will try to tell you that so-called "shared nothing" architectures scale well. While that may be literally true, I don't know of many interesting multi-user applications which can be implemented with no sharing of resources between different nodes of the cluster. What these silly people are really thinking of is a "share nothing except for the database" architecture. Of course, sharing the database is the primary problem with scaling a multi-user application—so the claim that this architecture is highly scalable is absurd, and tells you a lot about the kind of applications that these folks spend most of their time working on.

Almost anything we can possibly do to share the database *less often* is worth doing.

This calls for a cache. Well, not just one cache. A well designed Seam application will feature a rich, multi-layered caching strategy that impacts every layer of the application:

- The database, of course, has its own cache. This is super-important, but can't scale like a cache in the application tier.
- Your ORM solution (Hibernate, or some other JPA implementation) has a second-level cache of data from the database. This is a very powerful capability, but is often misused. In a clustered environment, keeping the data in the cache transactionally consistent across the whole cluster, and with the database, is quite expensive. It makes most sense for data which is shared between many users, and is updated rarely. In traditional stateless architectures, people often try to use the second-level cache for conversational state. This is always bad, and is especially wrong in Seam.
- The Seam conversation context is a cache of conversational state. Components you put into the conversation context can hold and cache state relating to the current user interaction.
- In particular, the Seam-managed persistence context (or an extended EJB container-managed persistence context associated with a conversation-scoped stateful session bean) acts as a cache of data that has been read in the current conversation. This cache tends to have a pretty high hitrate! Seam optimizes the replication of Seam-managed persistence contexts in a clustered environment, and there is no requirement for transactional consistency with the database (optimistic locking is sufficient) so you don't need to worry too much about the performance implications of this cache, unless you read thousands of objects into a single persistence context.
- The application can cache non-transactional state in the Seam application context. State kept in the application context is of course not visible to other nodes in the cluster.
- The application can cache transactional state using the Seam `pojoCache` component, which integrates JBossCache into the Seam environment. This state will be visible to other nodes if you run JBoss cache in a clustered mode.
- Finally, Seam lets you cache rendered fragments of a JSF page. Unlike the ORM second-level cache, this cache is not automatically invalidated when data changes, so you need to write application code to perform explicit invalidation, or set appropriate expiration policies.

For more information about the second-level cache, you'll need to refer to the documentation of your ORM solution, since this is an extremely complex topic. In this section we'll discuss the use of JBossCache directly, via the `pojoCache` component, or as the page fragment cache, via the `<s:cache>` control.

19.1. Using JBossCache in Seam

The built-in `pojoCache` component manages an instance of `org.jboss.cache.aop.PojoCache`. You can safely put any immutable Java object in the cache, and it will be replicated across the cluster (assuming that replication is enabled). If you want to keep mutable objects in the cache, you'll need to run the JBossCache bytecode preprocessor to ensure that changes to the objects will be automatically detected and replicated.

To use `pojoCache`, all you need to do is put the JBossCache jars in the classpath, and provide a resource named `treecache.xml` with an appropriate cache configuration. JBossCache has many scary and confusing configuration settings, so we won't discuss them here. Please refer to the JBossCache documentation for more information.

You can find a sample `treecache.xml` in `examples/blog/resources/treecache.xml`.

For an EAR deployment of Seam, we recommend that the JBossCache jars and configuration go directly into the EAR. Make sure you place both `jboss-cache.jar` and `jgroups.jar` in your EAR's `lib` folder.

Now you can inject the cache into any Seam component:

```
@Name("chatroom")
public class Chatroom {
    @In PojoCache.pojoCache;

    public void join(String username) {
        try
        {
            Set<String> userList = (Set<String>).pojoCache.get("chatroom", "userList");
            if (userList==null)
            {
                userList = new HashSet<String>();
               .pojoCache.put("chatroom", "userList", userList);
            }
            userList.put(username);
        }
        catch (CacheException ce)
        {
            throw new RuntimeException(ce);
        }
    }
}
```

If you want to have multiple JBossCache configurations in your application, use `components.xml`:

```
<core:pojo-cache name="myCache" cfg-resource-name="myown/cache.xml"/>
```

19.2. Page fragment caching

The most interesting user of JBossCache is the `<s:cache>` tag, Seam's solution to the problem of page fragment caching in JSF. `<s:cache>` uses `pojoCache` internally, so you need to follow the steps listed above before you can use it. (Put the jars in the EAR, wade through the scary configuration options, etc.)

`<s:cache>` is used for caching some rendered content which changes rarely. For example, the welcome page of our blog displays the recent blog entries:

```
<s:cache key="recentEntries-#{blog.id}" region="welcomePageFragments">
  <h:dataTable value="#{blog.recentEntries}" var="blogEntry">
    <h:column>
      <h3>#{blogEntry.title}</h3>
```

```
<div>
  <s:formattedText value="#{blogEntry.body}"/>
</div>
</h:column>
</h:dataTable>
</s:cache>
```

The key let's you have multiple cached versions of each page fragment. In this case, there is one cached version per blog. The region determines the JBossCache node that all version will be stored in. Different nodes may have different expiry policies. (That's the stuff you set up using the aforementioned scary configuration options.)

Of course, the big problem with `<s:cache>` is that it is too stupid to know when the underlying data changes (for example, when the blogger posts a new entry). So you need to evict the cached fragment manually:

```
public void post() {
    ...
    entityManager.persist(blogEntry);
    pojoCache.remove("welcomePageFragments", "recentEntries-" + blog.getId() );
}
```

Alternatively, if it is not critical that changes are immediately visible to the user, you could set a short expiry time on the JbossCache node.

Chapter 20. Web Services

Seam integrates with JBossWS to allow standard JEE web services to take full advantage of Seam's contextual framework, including support for conversational web services. This chapter walks through the steps required to allow web services to run within a Seam environment.

20.1. Configuration and Packaging

To allow Seam to intercept web service requests so that the necessary Seam contexts can be created for the request, a special SOAP handler must be configured; `org.jboss.seam.webservice.SOAPRequestHandler` is a `SOAPHandler` implementation that does the work of managing Seam's lifecycle during the scope of a web service request.

A special configuration file, `standard-jaxws-endpoint-config.xml` should be placed into the `META-INF` directory of the jar file that contains the web service classes. This file contains the following SOAP handler configuration:

```
<jaxws-config xmlns="urn:jboss:jaxws-config:2.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:javaee="http://java.sun.com/xml/ns/javaee"
  xsi:schemaLocation="urn:jboss:jaxws-config:2.0 jaxws-config_2_0.xsd">
  <endpoint-config>
    <config-name>Seam WebService Endpoint</config-name>
    <pre-handler-chains>
      <javaee:handler-chain>
        <javaee:protocol-bindings>##SOAP11_HTTP</javaee:protocol-bindings>
        <javaee:handler>
          <javaee:handler-name>SOAP Request Handler</javaee:handler-name>
          <javaee:handler-class>org.jboss.seam.webservice.SOAPRequestHandler</javaee:handler-class>
        </javaee:handler>
      </javaee:handler-chain>
    </pre-handler-chains>
  </endpoint-config>
</jaxws-config>
```

20.2. Conversational Web Services

So how are conversations propagated between web service requests? Seam uses a SOAP header element present in both the SOAP request and response messages to carry the conversation ID from the consumer to the service, and back again. Here's an example of a web service request that contains a conversation ID:

```
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:seam="http://seambay.example.seam.jboss.org/">
  <soapenv:Header>
    <seam:conversationId xmlns:seam='http://www.jboss.org/seam/webservice'>2</seam:conversationId>
  </soapenv:Header>
  <soapenv:Body>
    <seam:confirmAuction/>
  </soapenv:Body>
</soapenv:Envelope>
```

As you can see in the above SOAP message, there is a `conversationId` element within the SOAP header that contains the conversation ID for the request, in this case 2. Unfortunately, because web services may be consumed by a variety of web service clients written in a variety of languages, it is up to the developer to implement conversation ID propagation between individual web services that are intended to be used within the

scope of a single conversation.


An important thing to note is that the `conversationId` header element must be qualified with a namespace of `http://www.jboss.org/seam/webservice`, otherwise Seam will not be able to read the conversation ID from the request. Here's an example of a response to the above request message:

```
<env:Envelope xmlns:env='http://schemas.xmlsoap.org/soap/envelope/'>
  <env:Header>
    <seam:conversationId xmlns:seam='http://www.jboss.org/seam/webservice'>2</seam:conversationId>
  </env:Header>
  <env:Body>
    <confirmAuctionResponse xmlns="http://seambay.example.seam.jboss.org/">
  </env:Body>
</env:Envelope>
```

As you can see, the response message contains the same `conversationId` element as the request.

20.2.1. A Recommended Strategy

As web services must be implemented as either a stateless session bean or POJO, it is recommended that for conversational web services, the web service acts as a facade to a conversational Seam component.



If the web service is written as a stateless session bean, then it is also possible to make it a Seam component by giving it a `@Name`. Doing this allows Seam's bijection (and other) features to be used in the web service class itself.

20.3. An example web service

Let's walk through an example web service. The code in this section all comes from the `seamBay` example application in Seam's `/examples` directory, and follows the recommended strategy as described in the previous section. Let's first take a look at the web service class and one of its web service methods:

```
@Stateless
@WebService(name = "AuctionService", serviceName = "AuctionService")
public class AuctionService implements AuctionServiceRemote
{
  @WebMethod
```

```
public boolean login(String username, String password)
{
    Identity.instance().setUsername(username);
    Identity.instance().setPassword(password);
    Identity.instance().login();
    return Identity.instance().isLoggedIn();
}

// snip
}
```

As you can see, our web service is a stateless session bean, and is annotated using the JWS annotations from the `javax.jws` package, as defined by JSR-181. The `@WebService` annotation tells the container that this class implements a web service, and the `@WebMethod` annotation on the `login()` method identifies the method as a web service method. The `name` and `serviceName` attributes in the `@WebService` annotation are optional.

As is required by the specification, each method that is to be exposed as a web service method must also be declared in the remote interface of the web service class (when the web service is a stateless session bean). In the above example, the `AuctionServiceRemote` interface must declare the `login()` method as it is annotated as a `@WebMethod`.

As you can see in the above code, the web service implements a `login()` method that delegates to Seam's built-in `Identity` component. In keeping with our recommended strategy, the web service is written as a simple facade, passing off the real work to a Seam component. This allows for the greatest reuse of business logic between web services and other clients.

Let's look at another example. This web service method begins a new conversation by delegating to the `AuctionAction.createAuction()` method:

```
@WebMethod
public void createAuction(String title, String description, int categoryId)
{
    AuctionAction action = (AuctionAction) Component.getInstance(AuctionAction.class, true);
    action.createAuction();
    action.setDetails(title, description, categoryId);
}
```

And here's the code from `AuctionAction`:

```
@Begin
public void createAuction()
{
    auction = new Auction();
    auction.setAccount(authenticatedAccount);
    auction.setStatus(Auction.STATUS_UNLISTED);
    durationDays = DEFAULT_AUCTION_DURATION;
}
```

From this we can see how web services can participate in long running conversations, by acting as a facade and delegating the real work to a conversational Seam component.

Chapter 21. Remoting

Seam provides a convenient method of remotely accessing components from a web page, using AJAX (Asynchronous Javascript and XML). The framework for this functionality is provided with almost no up-front development effort - your components only require simple annotating to become accessible via AJAX. This chapter describes the steps required to build an AJAX-enabled web page, then goes on to explain the features of the Seam Remoting framework in more detail.

21.1. Configuration

To use remoting, the Seam Resource servlet must first be configured in your `web.xml` file:

```
<servlet>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <servlet-class>org.jboss.seam.servlet.SeamResourceServlet</servlet-class>
</servlet>

<servlet-mapping>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <url-pattern>/seam/resource/*</url-pattern>
</servlet-mapping>
```

The next step is to import the necessary Javascript into your web page. There are a minimum of two scripts that must be imported. The first one contains all the client-side framework code that enables remoting functionality:

```
<script type="text/javascript" src="seam/resource/remoting/resource/remote.js"></script>
```

The second script contains the stubs and type definitions for the components you wish to call. It is generated dynamically based on the local interface of your components, and includes type definitions for all of the classes that can be used to call the remotable methods of the interface. The name of the script reflects the name of your component. For example, if you have a stateless session bean annotated with `@Name("customerAction")`, then your script tag should look like this:

```
<script type="text/javascript"
      src="seam/resource/remoting/interface.js?customerAction"></script>
```

If you wish to access more than one component from the same page, then include them all as parameters of your script tag:

```
<script type="text/javascript"
      src="seam/resource/remoting/interface.js?customerAction&accountAction"></script>
```

Alternatively, you may use the `s:remote` tag to import the required Javascript. Separate each component or class name you wish to import with a comma:

```
<s:remote include="customerAction,accountAction"/>
```

21.2. The "Seam" object

Client-side interaction with your components is all performed via the Seam Javascript object. This object is defined in `remote.js`, and you'll be using it to make asynchronous calls against your component. It is split into

two areas of functionality; `Seam.Component` contains methods for working with components and `Seam.Remoting` contains methods for executing remote requests. The easiest way to become familiar with this object is to start with a simple example.

21.2.1. A Hello World example

Let's step through a simple example to see how the Seam object works. First of all, let's create a new Seam component called `helloAction`.

```
@Stateless
@Name("helloAction")
public class HelloAction implements HelloLocal {
    public String sayHello(String name) {
        return "Hello, " + name;
    }
}
```

You also need to create a local interface for our new component - take special note of the `@WebRemote` annotation, as it's required to make our method accessible via remoting:

```
@Local
public interface HelloLocal {
    @WebRemote
    public String sayHello(String name);
}
```

That's all the server-side code we need to write. Now for our web page - create a new page and import the `helloAction` component:

```
<s:remote include="helloAction"/>
```

To make this a fully interactive user experience, let's add a button to our page:

```
<button onclick="javascript:sayHello()">Say Hello</button>
```

We'll also need to add some more script to make our button actually do something when it's clicked:

```
<script type="text/javascript">
    /*

    function sayHello() {
        var name = prompt("What is your name?");
        Seam.Component.getInstance("helloAction").sayHello(name, sayHelloCallback);
    }

    function sayHelloCallback(result) {
        alert(result);
    }

    // ]]&gt;
&lt;/script&gt;</pre>
</div>
<div data-bbox="81 815 925 884" data-label="Text">
<p>We're done! Deploy your application and browse to your page. Click the button, and enter a name when prompted. A message box will display the hello message confirming that the call was successful. If you want to save some time, you'll find the full source code for this Hello World example in Seam's <code>/examples/remoting/helloworld</code> directory.</p>
</div>
<div data-bbox="81 896 925 931" data-label="Text">
<p>So what does the code of our script actually do? Let's break it down into smaller pieces. To start with, you can see from the Javascript code listing that we have implemented two methods - the first method is responsible for</p>
</div>
<div data-bbox="81 962 254 979" data-label="Page-Footer">JBoss Seam 2.0.1.CR1</div>
<div data-bbox="883 962 925 979" data-label="Page-Footer">205</div>
```

prompting the user for their name and then making a remote request. Take a look at the following line:

```
Seam.Component.getInstance("helloAction").sayHello(name, sayHelloCallback);
```

The first section of this line, `Seam.Component.getInstance("helloAction")` returns a proxy, or "stub" for our `helloAction` component. We can invoke the methods of our component against this stub, which is exactly what happens with the remainder of the line: `sayHello(name, sayHelloCallback);`.

What this line of code in its completeness does, is invoke the `sayHello` method of our component, passing in `name` as a parameter. The second parameter, `sayHelloCallback` isn't a parameter of our component's `sayHello` method, instead it tells the Seam Remoting framework that once it receives the response to our request, it should pass it to the `sayHelloCallback` Javascript method. This callback parameter is entirely optional, so feel free to leave it out if you're calling a method with a `void` return type or if you don't care about the result.

The `sayHelloCallback` method, once receiving the response to our remote request then pops up an alert message displaying the result of our method call.

21.2.2. Seam.Component

The `Seam.Component` Javascript object provides a number of client-side methods for working with your Seam components. The two main methods, `newInstance()` and `getInstance()` are documented in the following sections however their main difference is that `newInstance()` will always create a new instance of a component type, and `getInstance()` will return a singleton instance.

Seam.Component.newInstance()

Use this method to create a new instance of an entity or Javabeen component. The object returned by this method will have the same getter/setter methods as its server-side counterpart, or alternatively if you wish you can access its fields directly. Take the following Seam entity component for example:

```
@Name("customer")
@Entity
public class Customer implements Serializable
{
    private Integer customerId;
    private String firstName;
    private String lastName;

    @Column public Integer getCustomerId() {
        return customerId;
    }

    public void setCustomerId(Integer customerId) {
        this.customerId = customerId;
    }

    @Column public String getFirstName() {
        return firstName;
    }

    public void setFirstName(String firstName) {
        this.firstName = firstName;
    }

    @Column public String getLastName() {
        return lastName;
    }

    public void setLastName(String lastName) {
        this.lastName = lastName;
    }
}
```

```
}
```

To create a client-side Customer you would write the following code:

```
var customer = Seam.Component.newInstance("customer");
```

Then from here you can set the fields of the customer object:

```
customer.setFirstName("John");  
// Or you can set the fields directly  
customer.lastName = "Smith";
```

Seam.Component.getInstance()

The `getInstance()` method is used to get a reference to a Seam session bean component stub, which can then be used to remotely execute methods against your component. This method returns a singleton for the specified component, so calling it twice in a row with the same component name will return the same instance of the component.

To continue our example from before, if we have created a new customer and we now wish to save it, we would pass it to the `saveCustomer()` method of our `customerAction` component:

```
Seam.Component.getInstance("customerAction").saveCustomer(customer);
```

Seam.Component.getComponentName()

Passing an object into this method will return its component name if it is a component, or `null` if it is not.

```
if (Seam.Component.getComponentName(instance) == "customer")  
    alert("Customer");  
else if (Seam.Component.getComponentName(instance) == "staff")  
    alert("Staff member");
```

21.2.3. Seam.Remoting

Most of the client side functionality for Seam Remoting is contained within the `Seam.Remoting` object. While you shouldn't need to directly call most of its methods, there are a couple of important ones worth mentioning.

Seam.Remoting.createType()

If your application contains or uses Javabeans classes that aren't Seam components, you may need to create these types on the client side to pass as parameters into your component method. Use the `createType()` method to create an instance of your type. Pass in the fully qualified Java class name as a parameter:

```
var widget = Seam.Remoting.createType("com.acme.widgets.MyWidget");
```

Seam.Remoting.getTypeName()

This method is the equivalent of `Seam.Component.getComponentName()` but for non-component types. It will return the name of the type for an object instance, or `null` if the type is not known. The name is the fully qualified name of the type's Java class.

21.3. Evaluating EL Expressions

Seam Remoting also supports the evaluation of EL expressions, which provides another convenient method for retrieving data from the server. Using the `Seam.Remoting.eval()` function, an EL expression can be remotely evaluated on the server and the resulting value returned to a client-side callback method. This function accepts two parameters, the first being the EL expression to evaluate, and the second being the callback method to invoke with the value of the expression. Here's an example:

```
function customersCallback(customers) {
    for (var i = 0; i < customers.length; i++) {
        alert("Got customer: " + customers[i].getName());
    }
}

Seam.Remoting.eval("#{customers}", customersCallback);
```

In this example, the expression `#{customers}` is evaluated by Seam, and the value of the expression (in this case a list of `Customer` objects) is returned to the `customersCallback()` method. It is important to remember that the objects returned this way must have their types imported (via `s:remote`) to be able to work with them in Javascript. So to work with a list of customer objects, it is required to import the customer type:

```
<s:remote include="customer"/>
```

21.4. Client Interfaces

In the configuration section above, the interface, or "stub" for our component is imported into our page either via `seam/resource/remoting/interface.js`:

or using the `s:remote` tag:

```
<script type="text/javascript"
    src="seam/resource/remoting/interface.js?customerAction"></script>
```

```
<s:remote include="customerAction"/>
```

By including this script in our page, the interface definitions for our component, plus any other components or types that are required to execute the methods of our component are generated and made available for the remoting framework to use.

There are two types of client stub that can be generated, "executable" stubs and "type" stubs. Executable stubs are behavioural, and are used to execute methods against your session bean components, while type stubs contain state and represent the types that can be passed in as parameters or returned as a result.

The type of client stub that is generated depends on the type of your Seam component. If the component is a session bean, then an executable stub will be generated, otherwise if it's an entity or `JavaBean`, then a type stub will be generated. There is one exception to this rule; if your component is a `JavaBean` (ie it is not a session bean nor an entity bean) and any of its methods are annotated with `@WebRemote`, then an executable stub will be generated for it instead of a type stub. This allows you to use remoting to call methods of your `JavaBean` components in a non-EJB environment where you don't have access to session beans.

21.5. The Context

The Seam Remoting Context contains additional information which is sent and received as part of a remoting request/response cycle. At this stage it only contains the conversation ID but may be expanded in the future.

21.5.1. Setting and reading the Conversation ID

If you intend on using remote calls within the scope of a conversation then you need to be able to read or set the conversation ID in the Seam Remoting Context. To read the conversation ID after making a remote request call `Seam.Remoting.getContext().getConversationId()`. To set the conversation ID before making a request, call `Seam.Remoting.getContext().setConversationId()`.

If the conversation ID hasn't been explicitly set with `Seam.Remoting.getContext().setConversationId()`, then it will be automatically assigned the first valid conversation ID that is returned by any remoting call. If you are working with multiple conversations within your page, then you may need to explicitly set the conversation ID before each call. If you are working with just a single conversation, then you don't need to do anything special.

21.5.2. Remote calls within the current conversation scope

In some circumstances it may be required to make a remote call within the scope of the current view's conversation. To do this, you must explicitly set the conversation ID to that of the view before making the remote call. This small snippet of JavaScript will set the conversation ID that is used for remoting calls to the current view's conversation ID:

```
Seam.Remoting.getContext().setConversationId( #{conversation.id} );
```

21.6. Batch Requests

Seam Remoting allows multiple component calls to be executed within a single request. It is recommended that this feature is used wherever it is appropriate to reduce network traffic.

The method `Seam.Remoting.startBatch()` will start a new batch, and any component calls executed after starting a batch are queued, rather than being sent immediately. When all the desired component calls have been added to the batch, the `Seam.Remoting.executeBatch()` method will send a single request containing all of the queued calls to the server, where they will be executed in order. After the calls have been executed, a single response containing all return values will be returned to the client and the callback functions (if provided) triggered in the same order as execution.

If you start a new batch via the `startBatch()` method but then decide you don't want to send it, the `Seam.Remoting.cancelBatch()` method will discard any calls that were queued and exit the batch mode.

To see an example of a batch being used, take a look at `/examples/remoting/chatroom`.

21.7. Working with Data types

21.7.1. Primitives / Basic Types

This section describes the support for basic data types. On the server side these values are generally compatible with either their primitive type or their corresponding wrapper class.

String

Simply use Javascript String objects when setting String parameter values.

Number

There is support for all number types supported by Java. On the client side, number values are always serialized as their String representation and then on the server side they are converted to the correct destination type. Conversion into either a primitive or wrapper type is supported for Byte, Double, Float, Integer, Long and Short types.

Boolean

Booleans are represented client side by Javascript Boolean values, and server side by a Java boolean.

21.7.2. JavaBeans

In general these will be either Seam entity or JavaBean components, or some other non-component class. Use the appropriate method (either `Seam.Component.newInstance()` for Seam components or `Seam.Remoting.createType()` for everything else) to create a new instance of the object.

It is important to note that only objects that are created by either of these two methods should be used as parameter values, where the parameter is not one of the other valid types mentioned anywhere else in this section. In some situations you may have a component method where the exact parameter type cannot be determined, such as:

```
@Name("myAction")
public class MyAction implements MyActionLocal {
    public void doSomethingWithObject(Object obj) {
        // code
    }
}
```

In this case you might want to pass in an instance of your myWidget component, however the interface for myAction won't include myWidget as it is not directly referenced by any of its methods. To get around this, MyWidget needs to be explicitly imported:

```
<s:remote include="myAction,myWidget"/>
```

This will then allow a myWidget object to be created with `Seam.Component.newInstance("myWidget")`, which can then be passed to `myAction.doSomethingWithObject()`.

21.7.3. Dates and Times

Date values are serialized into a String representation that is accurate to the millisecond. On the client side, use a Javascript Date object to work with date values. On the server side, use any `java.util.Date` (or descendent, such as `java.sql.Date` or `java.sql.Timestamp` class.

21.7.4. Enums

On the client side, enums are treated the same as Strings. When setting the value for an enum parameter, simply use the String representation of the enum. Take the following component as an example:

```
@Name("paintAction")
public class paintAction implements paintLocal {
    public enum Color {red, green, blue, yellow, orange, purple};

    public void paint(Color color) {
        // code
    }
}
```

To call the `paint()` method with the color red, pass the parameter value as a String literal:

```
Seam.Component.getInstance("paintAction").paint("red");
```

The inverse is also true - that is, if a component method returns an enum parameter (or contains an enum field anywhere in the returned object graph) then on the client-side it will be represented as a String.

21.7.5. Collections

Bags

Bags cover all collection types including arrays, collections, lists, sets, (but excluding Maps - see the next section for those), and are implemented client-side as a Javascript array. When calling a component method that accepts one of these types as a parameter, your parameter should be a Javascript array. If a component method returns one of these types, then the return value will also be a Javascript array. The remoting framework is clever enough on the server side to convert the bag to an appropriate type for the component method call.

Maps

As there is no native support for Maps within Javascript, a simple Map implementation is provided with the Seam Remoting framework. To create a Map which can be used as a parameter to a remote call, create a new `Seam.Remoting.Map` object:

```
var map = new Seam.Remoting.Map();
```

This Javascript implementation provides basic methods for working with Maps: `size()`, `isEmpty()`, `keySet()`, `values()`, `get(key)`, `put(key, value)`, `remove(key)` and `contains(key)`. Each of these methods are equivalent to their Java counterpart. Where the method returns a collection, such as `keySet()` and `values()`, a Javascript Array object will be returned that contains the key or value objects (respectively).

21.8. Debugging

To aid in tracking down bugs, it is possible to enable a debug mode which will display the contents of all the packets send back and forth between the client and server in a popup window. To enable debug mode, either execute the `setDebug()` method in Javascript:

```
Seam.Remoting.setDebug(true);
```

Or configure it via `components.xml`:

```
<remoting:remoting debug="true"/>
```

To turn off debugging, call `setDebug(false)`. If you want to write your own messages to the debug log, call `Seam.Remoting.log(message)`.

21.9. The Loading Message

The default loading message that appears in the top right corner of the screen can be modified, its rendering customised or even turned off completely.

21.9.1. Changing the message

To change the message from the default "Please Wait..." to something different, set the value of `Seam.Remoting.loadingMessage`:

```
Seam.Remoting.loadingMessage = "Loading...";
```

21.9.2. Hiding the loading message

To completely suppress the display of the loading message, override the implementation of `displayLoadingMessage()` and `hideLoadingMessage()` with functions that instead do nothing:

```
// don't display the loading indicator
Seam.Remoting.displayLoadingMessage = function() {};
Seam.Remoting.hideLoadingMessage = function() {};
```

21.9.3. A Custom Loading Indicator

It is also possible to override the loading indicator to display an animated icon, or anything else that you want. To do this override the `displayLoadingMessage()` and `hideLoadingMessage()` messages with your own implementation:

```
Seam.Remoting.displayLoadingMessage = function() {
    // Write code here to display the indicator
};

Seam.Remoting.hideLoadingMessage = function() {
    // Write code here to hide the indicator
};
```

21.10. Controlling what data is returned

When a remote method is executed, the result is serialized into an XML response that is returned to the client. This response is then unmarshaled by the client into a Javascript object. For complex types (i.e. Javabeans) that include references to other objects, all of these referenced objects are also serialized as part of the response. These objects may reference other objects, which may reference other objects, and so forth. If left unchecked, this object "graph" could potentially be enormous, depending on what relationships exist between your objects. And as a side issue (besides the potential verbosity of the response), you might also wish to prevent sensitive information from being exposed to the client.

Seam Remoting provides a simple means to "constrain" the object graph, by specifying the `exclude` field of the remote method's `@WebRemote` annotation. This field accepts a String array containing one or more paths specified using dot notation. When invoking a remote method, the objects in the result's object graph that match these paths are excluded from the serialized result packet.

For all our examples, we'll use the following widget class:

```
@Name("widget")
public class Widget
{
    private String value;
    private String secret;
    private Widget child;
    private Map<String,Widget> widgetMap;
    private List<Widget> widgetList;

    // getters and setters for all fields
}
```

21.10.1. Constraining normal fields

If your remote method returns an instance of `Widget`, but you don't want to expose the `secret` field because it contains sensitive information, you would constrain it like this:

```
@WebRemote(exclude = {"secret"})
public Widget getWidget();
```

The value `"secret"` refers to the `secret` field of the returned object. Now, suppose that we don't care about exposing this particular field to the client. Instead, notice that the `Widget` value that is returned has a field `child` that is also a `Widget`. What if we want to hide the `child's` `secret` value instead? We can do this by using dot notation to specify this field's path within the result's object graph:

```
@WebRemote(exclude = {"child.secret"})
public Widget getWidget();
```

21.10.2. Constraining Maps and Collections

The other place that objects can exist within an object graph are within a `Map` or some kind of collection (`List`, `Set`, `Array`, etc). Collections are easy, and are treated like any other field. For example, if our `Widget` contained a list of other `Widgets` in its `widgetList` field, to constrain the `secret` field of the `Widgets` in this list the annotation would look like this:

```
@WebRemote(exclude = {"widgetList.secret"})
public Widget getWidget();
```

To constrain a `Map's` key or value, the notation is slightly different. Appending `[key]` after the `Map's` field name will constrain the `Map's` key object values, while `[value]` will constrain the value object values. The following example demonstrates how the values of the `widgetMap` field have their `secret` field constrained:

```
@WebRemote(exclude = {"widgetMap[value].secret"})
public Widget getWidget();
```

21.10.3. Constraining objects of a specific type

There is one last notation that can be used to constrain the fields of a type of object no matter where in the result's object graph it appears. This notation uses either the name of the component (if the object is a Seam component) or the fully qualified class name (only if the object is not a Seam component) and is expressed using square brackets:

```
@WebRemote(exclude = {"[widget].secret"})
public Widget getWidget();
```

21.10.4. Combining Constraints

Constraints can also be combined, to filter objects from multiple paths within the object graph:

```
@WebRemote(exclude = {"widgetList.secret", "widgetMap[value].secret"})
public Widget getWidget();
```

21.11. JMS Messaging

Seam Remoting provides experimental support for JMS Messaging. This section describes the JMS support that is currently implemented, but please note that this may change in the future. It is currently not recommended that this feature is used within a production environment.

21.11.1. Configuration

Before you can subscribe to a JMS topic, you must first configure a list of the topics that can be subscribed to by Seam Remoting. List the topics under `org.jboss.seam.remoting.messaging.subscriptionRegistry.allowedTopics` in `seam.properties`, `web.xml` or `components.xml`.

```
<remoting:remoting poll-timeout="5" poll-interval="1"/>
```

21.11.2. Subscribing to a JMS Topic

The following example demonstrates how to subscribe to a JMS Topic:

```
function subscriptionCallback(message)
{
    if (message instanceof Seam.Remoting.TextMessage)
        alert("Received message: " + message.getText());
}

Seam.Remoting.subscribe("topicName", subscriptionCallback);
```

The `Seam.Remoting.subscribe()` method accepts two parameters, the first being the name of the JMS Topic to subscribe to, the second being the callback function to invoke when a message is received.

There are two types of messages supported, Text messages and Object messages. If you need to test for the type of message that is passed to your callback function you can use the `instanceof` operator to test whether the message is a `Seam.Remoting.TextMessage` or `Seam.Remoting.ObjectMessage`. A `TextMessage` contains the text value in its `text` field (or alternatively call `getText()` on it), while an `ObjectMessage` contains its object value in its `value` field (or call its `getValue()` method).

21.11.3. Unsubscribing from a Topic

To unsubscribe from a topic, call `Seam.Remoting.unsubscribe()` and pass in the topic name:

```
Seam.Remoting.unsubscribe("topicName");
```

21.11.4. Tuning the Polling Process

There are two parameters which you can modify to control how polling occurs. The first one is `Seam.Remoting.pollInterval`, which controls how long to wait between subsequent polls for new messages. This parameter is expressed in seconds, and its default setting is 10.

The second parameter is `Seam.Remoting.pollTimeout`, and is also expressed as seconds. It controls how long a request to the server should wait for a new message before timing out and sending an empty response. Its default is 0 seconds, which means that when the server is polled, if there are no messages ready for delivery then an empty response will be immediately returned.

Caution should be used when setting a high `pollTimeout` value; each request that has to wait for a message means that a server thread is tied up until a message is received, or until the request times out. If many such requests are being served simultaneously, it could mean a large number of threads become tied up because of this reason.

It is recommended that you set these options via `components.xml`, however they can be overridden via JavaScript if desired. The following example demonstrates how to configure the polling to occur much more aggressively. You should set these parameters to suitable values for your application:

Via `components.xml`:

```
<remoting:remoting poll-timeout="5" poll-interval="1"/>
```

Via JavaScript:

```
// Only wait 1 second between receiving a poll response and sending the next poll request.
Seam.Remoting.pollInterval = 1;

// Wait up to 5 seconds on the server for new messages
Seam.Remoting.pollTimeout = 5;
```

Chapter 22. Seam and the Google Web Toolkit

For those that prefer to use the Google Web Toolkit (GWT) to develop dynamic AJAX applications, Seam provides an integration layer that allows GWT widgets to interact directly with Seam components.

To use GWT, we assume that you are already familiar with the GWT tools - more information can be found at <http://code.google.com/webtoolkit/>. This chapter does not attempt to explain how GWT works or how to use it.

22.1. Configuration

There is no special configuration required to use GWT in a Seam application, however the Seam resource servlet must be installed. See Chapter 25, *Configuring Seam and packaging Seam applications* for details.

22.2. Preparing your component

The first step in preparing a Seam component to be called via GWT, is to create both synchronous and asynchronous service interfaces for the methods you wish to call. Both of these interfaces should extend the GWT interface `com.google.gwt.user.client.rpc.RemoteService`:

```
public interface MyService extends RemoteService
{
    public String askIt(String question);
}
```

The asynchronous interface should be identical, except that it also contains an additional `AsyncCallback` parameter for each of the methods it declares:

```
public interface MyServiceAsync extends RemoteService
{
    public void askIt(String question, AsyncCallback callback);
}
```

The asynchronous interface, in this example `MyServiceAsync`, will be implemented by GWT and should never be implemented directly.

The next step, is to create a Seam component that implements the synchronous interface:

```
@Name("org.jboss.seam.example.remoting.gwt.client.MyService")
public class ServiceImpl implements MyService
{
    @WebRemote
    public String askIt(String question)
    {
        if (!validate(question))
        {
            throw new IllegalStateException("Hey, this shouldn't happen, I checked on the client, " +
                "but its always good to double check.");
        }
        return "42. Its the real question that you seek now.";
    }

    public boolean validate(String q)
    {
        ValidationUtility util = new ValidationUtility();
        return util.isValid(q);
    }
}
```

The methods that should be made accessible via GWT need to be annotated with the `@WebRemote` annotation, which is required for all web-remoteable methods.

22.3. Hooking up a GWT widget to the Seam component

The next step, is to write a method that returns the asynchronous interface to the component. This method can be located inside the widget class, and will be used by the widget to obtain a reference to the asynchronous client stub:

```
private MyServiceAsync getService()
{
    String endpointURL = GWT.getModuleBaseURL() + "seam/resource/gwt";

    MyServiceAsync svc = (MyServiceAsync) GWT.create(MyService.class);
    ((ServiceDefTarget) svc).setServiceEntryPoint(endpointURL);
    return svc;
}
```

The final step is to write the widget code that invokes the method on the client stub. The following example creates a simple user interface with a label, text input and a button:

```
public class AskQuestionWidget extends Composite
{
    private AbsolutePanel panel = new AbsolutePanel();

    public AskQuestionWidget()
    {
        Label lbl = new Label("OK, what do you want to know?");
        panel.add(lbl);
        final TextBox box = new TextBox();
        box.setText("What is the meaning of life?");
        panel.add(box);
        Button ok = new Button("Ask");
        ok.addClickListener(new ClickListener()
        {
            public void onClick(Widget w)
            {
                ValidationUtility valid = new ValidationUtility();
                if (!valid.isValid(box.getText()))
                {
                    Window.alert("A question has to end with a '?'");
                }
                else
                {
                    askServer(box.getText());
                }
            }
        });
        panel.add(ok);

        initWidget(panel);
    }

    private void askServer(String text)
    {
        getService().askIt(text, new AsyncCallback()
        {
            public void onFailure(Throwable t)
            {
                Window.alert(t.getMessage());
            }

            public void onSuccess(Object data)
            {
            }
        });
    }
}
```

```

        Window.alert((String) data);
    });
}
...

```

When clicked, the button invokes the `askServer()` method passing the contents of the input text (in this example, validation is also performed to ensure that the input is a valid question). The `askServer()` method acquires a reference to the asynchronous client stub (returned by the `getService()` method) and invokes the `askIt()` method. The result (or error message if the call fails) is shown in an alert window.

HelloWorld

This is an example of a host page for the HelloWorld application. You can attach a Web Toolkit module to any HTML page you like, making it easy to add bits of AJAX functionality to existing pages without starting from scratch.

OK, what do you want to know?

The complete code for this example can be found in the Seam distribution in the `examples/remoting/gwt` directory.

22.4. GWT Ant Targets

For deployment of GWT apps, there is a compile-to-Javascript step (which compacts and obfuscates the code). There is an ant utility which can be used instead of the command line or GUI utility that GWT provides. To use this, you will need to have the ant task jar in your ant classpath, as well as GWT downloaded (which you will need for hosted mode anyway).

Then, in your ant file, place (near the top of your ant file):

```

<taskdef uri="antlib:de.samaflost.gwttasks"
         resource="de/samaflost/gwttasks/antlib.xml"
         classpath="./lib/gwttasks.jar"/>

<property file="build.properties"/>

```

Create a `build.properties` file, which has the contents:

```
gwt.home=/gwt_home_dir
```

This of course should point to the directory where GWT is installed. Then to use it, create a target:

```

<!-- the following are are handy utilities for doing GWT development.
     To use GWT, you will of course need to download GWT seperately -->
<target name="gwt-compile">
    <!-- in this case, we are "re homing" the gwt generated stuff, so in this case
         we can only have one GWT module - we are doing this deliberately to keep the URL short -->
    <delete>
        <fileset dir="view"/>
    </delete>
    <gwt:compile outDir="build/gwt"
                 gwtHome="${gwt.home}"
                 classBase="${gwt.module.name}"
                 sourceclasspath="src"/>
    <copy todir="view">
        <fileset dir="build/gwt/${gwt.module.name}"/>
    </copy>
</target>

```

```
</copy>  
</target>
```

This target when called will compile the GWT application, and copy it to the specified directory (which would be in the webapp part of your war - remember GWT generates HTML and Javascript artifacts). You never edit the resulting code that `gwt-compile` generates - you always edit in the GWT source directory.

Remember that GWT comes with a hosted mode browser - you should be using that if you are developing with GWT. If you aren't using that, and are just compiling it each time, you aren't getting the most out of the toolkit (in fact, if you can't or won't use the hosted mode browser, I would go far as to say you should NOT be using GWT at all - it's that valuable!).

Chapter 23. Spring Framework integration

The Spring integration module allows easy migration of Spring-based projects to Seam and allows Spring applications to take advantage of key Seam features like conversations and Seam's more sophisticated persistence context management.

Note! The Spring integration code is included in the `jboss-seam-ioc` library. This dependency is required for all seam-spring integration techniques covered in this chapter.

Seam's support for Spring provides the ability to:

- inject Seam component instances into Spring beans
- inject Spring beans into Seam components
- turn Spring beans into Seam components
- allow Spring beans to live in any Seam context
- start a spring `WebApplicationContext` with a Seam component
- Support for Spring `PlatformTransactionManagement`
- provides a Seam managed replacement for Spring's `OpenEntityManagerInViewFilter` and `OpenSessionInViewFilter`
- Support for Spring `TaskExecutors` to back `@Asynchronous` calls

23.1. Injecting Seam components into Spring beans

Injecting Seam component instances into Spring beans is accomplished using the `<seam:instance/>` namespace handler. To enable the Seam namespace handler, the Seam namespace must be added to the Spring beans definition file:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:seam="http://jboss.com/products/seam/spring-seam"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
    http://jboss.com/products/seam/spring-seam
    http://jboss.com/products/seam/spring-seam-2.0.xsd">
```

Now any Seam component may be injected into any Spring bean:

```
<bean id="someSpringBean" class="SomeSpringBeanClass" scope="prototype">
  <property name="someProperty">
    <seam:instance name="someComponent"/>
  </property>
</bean>
```

An EL expression may be used instead of a component name:

```
<bean id="someSpringBean" class="SomeSpringBeanClass" scope="prototype">
  <property name="someProperty">
    <seam:instance name="#{someExpression}"/>
  </property>
</bean>
```

```
</bean>
```

Seam component instances may even be made available for injection into Spring beans by a Spring bean id.

```
<seam:instance name="someComponent" id="someSeamComponentInstance"/>
<bean id="someSpringBean" class="SomeSpringBeanClass" scope="prototype">
  <property name="someProperty" ref="someSeamComponentInstance">
</bean>
```

Now for the caveat!

Seam was designed from the ground up to support a stateful component model with multiple contexts. Spring was not. Unlike Seam bijection, Spring injection does not occur at method invocation time. Instead, injection happens only when the Spring bean is instantiated. So the instance available when the bean is instantiated will be the same instance that the bean uses for the entire life of the bean. For example, if a Seam CONVERSATION-scoped component instance is directly injected into a singleton Spring bean, that singleton will hold a reference to the same instance long after the conversation is over! We call this problem *scope impedance*. Seam bijection ensures that scope impedance is maintained naturally as an invocation flows through the system. In Spring, we need to inject a proxy of the Seam component, and resolve the reference when the proxy is invoked.

The `<seam:instance/>` tag lets us automatically proxy the Seam component.

```
<seam:instance id="seamManagedEM" name="someManagedEMComponent" proxy="true"/>
<bean id="someSpringBean" class="SomeSpringBeanClass">
  <property name="entityManager" ref="seamManagedEM">
</bean>
```

This example shows one way to use a Seam-managed persistence context from a Spring bean. (For a more robust way to use Seam-managed persistence contexts as a replacement for the Spring `OpenEntityManagerInView` filter see section on Using a Seam Managed Persistence Context in Spring)

23.2. Injecting Spring beans into Seam components

It is even easier to inject Spring beans into Seam component instances. Actually, there are two possible approaches:

- inject a Spring bean using an EL expression
- make the Spring bean a Seam component

We'll discuss the second option in the next section. The easiest approach is to access the Spring beans via EL.

The Spring `DelegatingVariableResolver` is an integration point Spring provides for integrating Spring with JSF. This `variableResolver` makes all Spring beans available in EL by their bean id. You'll need to add the `DelegatingVariableResolver` to `faces-config.xml`:

```
<application>
  <variable-resolver>
    org.springframework.web.jsf.DelegatingVariableResolver
  </variable-resolver>
</application>
```

Then you can inject Spring beans using `@In`:

```
@In("#{bookingService}")
private BookingService bookingService;
```

The use of Spring beans in EL is not limited to injection. Spring beans may be used anywhere that EL expressions are used in Seam: process and pageflow definitions, working memory assertions, etc...

23.3. Making a Spring bean into a Seam component

The `<seam:component/>` namespace handler can be used to make any Spring bean a Seam component. Just place the `<seam:component/>` tag within the declaration of the bean that you wish to be a Seam component:

```
<bean id="someSpringBean" class="SomeSpringBeanClass" scope="prototype">
  <seam:component/>
</bean>
```

By default, `<seam:component/>` will create a STATELESS Seam component with class and name provided in the bean definition. Occasionally, such as when a `FactoryBean` is used, the class of the Spring bean may not be the class appearing in the bean definition. In such cases the `class` should be explicitly specified. A Seam component name may be explicitly specified in cases where there is potential for a naming conflict.

The `scope` attribute of `<seam:component/>` may be used if you wish the Spring bean to be managed in a particular Seam scope. The Spring bean must be scoped to `prototype` if the Seam scope specified is anything other than `STATELESS`. Pre-existing Spring beans usually have a fundamentally stateless character, so this attribute is not usually needed.

23.4. Seam-scoped Spring beans

The Seam integration package also lets you use Seam's contexts as Spring 2.0 style custom scopes. This lets you declare any Spring bean in any of Seam's contexts. However, note once again that Spring's component model was never architected to support statefulness, so please use this feature with great care. In particular, clustering of session or conversation scoped Spring beans is deeply problematic, and care must be taken when injecting a bean or component from a wider scope into a bean of a narrower scope.

By specifying `<seam:configure-scopes/>` once in a Spring bean factory configuration, all of the Seam scopes will be available to Spring beans as custom scopes. To associate a Spring bean with a particular Seam scope, specify the Seam scope in the `scope` attribute of the bean definition.

```
<!-- Only needs to be specified once per bean factory-->
<seam:configure-scopes/>

...

<bean id="someSpringBean" class="SomeSpringBeanClass" scope="seam.CONVERSATION"/>
```

The prefix of the scope name may be changed by specifying the `prefix` attribute in the `configure-scopes` definition. (The default prefix is `seam`.)

By default an instance of a Spring Component registered in this way is not automatically created when referenced using `@In`. To have an instance auto-created you must either specify `@In(create=true)` at the injection point to identify a specific bean to be auto created or you can use the `default-auto-create` attribute of `configure-scopes` to make all spring beans who use a seam scope auto created.

Seam-scoped Spring beans defined this way can be injected into other Spring beans without the use of

`<seam:instance/>`. However, care must be taken to ensure scope impedance is maintained. The normal approach used in Spring is to specify `<aop:scoped-proxy/>` in the bean definition. However, Seam-scoped Spring beans are *not* compatible with `<aop:scoped-proxy/>`. So if you need to inject a Seam-scoped Spring bean into a singleton, `<seam:instance/>` must be used:

```
<bean id="someSpringBean" class="SomeSpringBeanClass" scope="seam.CONVERSATION"/>

...

<bean id="someSingleton">
    <property name="someSeamScopedSpringBean">
        <seam:instance name="someSpringBean" proxy="true"/>
    </property>
</bean>
```

23.5. Using Spring PlatformTransactionManagement

Spring provides an extensible transaction management abstraction with support for many transaction APIs (JPA, Hibernate, JDO, and JTA) Spring also provides tight integrations with many application server TransactionManagers such as Websphere and Weblogic. Spring transaction management exposes support for many advanced features such as nested transactions and supports full Java EE transaction propagation rules like `REQUIRES_NEW` and `NOT_SUPPORTED`. For more information see the spring documentation here [<http://static.springframework.org/spring/docs/2.0.x/reference/transaction.html>].

To configure Seam to use Spring transactions enable the SpringTransaction component like so:

```
<spring:spring-transaction platform-transaction-manager="#{transactionManager}" />
```

The `spring:spring-transaction` component will utilize Springs transaction synchronization capabilities for synchronization callbacks.

23.6. Using a Seam Managed Persistence Context in Spring

One of the most powerful features of Seam is its conversation scope and the ability to have an `EntityManager` open for the life of a conversation. This eliminates many of the problems associated with the detachment and re-attachment of entities as well as mitigates occurrences of the dreaded `LazyInitializationException`. Spring does not provide a way to manage an persistence context beyond the scope of a single web request (`OpenEntityManagerInViewFilter`). So, it would be nice if Spring developers could have access to a Seam managed persistence context using all of the same tools Spring provides for integration with JPA (e.g. `PersistenceAnnotationBeanPostProcessor`, `JpaTemplate`, etc.)

Seam provides a way for Spring to access a Seam managed persistence context with Spring's provided JPA tools bringing conversation scoped persistence context capabilities to Spring applications.

This integration work provides the following functionality:

- transparent access to a Seam managed persistence context using Spring provided tools
- access to Seam conversation scoped persistence contexts in a non web request (e.g. asynchronous quartz job)
- allows for using Seam managed persistence contexts with Spring managed transactions (will need to flush

the persistence context manually)

Spring's persistence context propagation model allows only one open EntityManager per EntityManagerFactory so the Seam integration works by wrapping an EntityManagerFactory around a Seam managed persistence context.

```
<bean id="seamEntityManagerFactory" class="org.jboss.seam.ioc.spring.SeamManagedEntityManagerFactoryB
    <property name="persistenceContextName" value="entityManager"/>
</bean>
```

Where 'persistenceContextName' is the name of the Seam managed persistence context component. By default this EntityManagerFactory has a unitName equal to the Seam component name or in this case 'entityManager'. If you wish to provide a different unitName you can do so by providing a persistenceUnitName like so:

```
<bean id="seamEntityManagerFactory" class="org.jboss.seam.ioc.spring.SeamManagedEntityManagerFactoryB
    <property name="persistenceContextName" value="entityManager"/>
    <property name="persistenceUnitName" value="bookingDatabase:extended"/>
</bean>
```

This EntityManagerFactory can then be used in any Spring provided tools. For example, using Spring's PersistenceAnnotationBeanPostProcessor is the exact same as before.

```
<bean class="org.springframework.orm.jpa.support.PersistenceAnnotationBeanPostProcessor"/>
```

If you define your real EntityManagerFactory in Spring but wish to use a Seam managed persistence context you can tell the PersistenceAnnotationBeanPostProcessor which persistenceUnitName you wish to use by default by specifying the defaultPersistenceUnitName property.

The applicationContext.xml might look like:

```
<bean id="entityManagerFactory" class="org.springframework.orm.jpa.LocalEntityManagerFactoryBean">
    <property name="persistenceUnitName" value="bookingDatabase"/>
</bean>
<bean id="seamEntityManagerFactory" class="org.jboss.seam.ioc.spring.SeamManagedEntityManagerFactoryB
    <property name="persistenceContextName" value="entityManager"/>
    <property name="persistenceUnitName" value="bookingDatabase:extended"/>
</bean>
<bean class="org.springframework.orm.jpa.support.PersistenceAnnotationBeanPostProcessor">
    <property name="defaultPersistenceUnitName" value="bookingDatabase:extended"/>
</bean>
```

The component.xml might look like:

```
<persistence:managed-persistence-context name="entityManager"
    auto-create="true" entity-manager-factory="#{entityManagerFactory}"/>
```

JpaTemplate and JpaDaoSupport are configured the same way for a Seam managed persistence context as they would be for a Seam managed persistence context.

```
<bean id="bookingService" class="org.jboss.seam.example.spring.BookingService">
    <property name="entityManagerFactory" ref="seamEntityManagerFactory"/>
</bean>
```

23.7. Using a Seam Managed Hibernate Session in Spring

The Seam Spring integration also provides support for complete access to a Seam managed Hibernate session

using spring's tools. This integration is very similar to the JPA integration.

Like Spring's JPA integration spring's propagation model allows only one open EntityManager per EntityManagerFactory per transaction??? to be available to spring tools. So, the Seam Session integration works by wrapping a proxy SessionFactory around a Seam managed Hibernate session context.

```
<bean id="seamSessionFactory" class="org.jboss.seam.ioc.spring.SeamManagedSessionFactoryBean">
    <property name="sessionName" value="hibernateSession"/>
</bean>
```

Where 'sessionName' is the name of the persistence:managed-hibernate-session component. This SessionFactory can then be used in any Spring provided tools. The integration also provides support for calls to SessionFactory.getCurrentInstance() as long as you call getCurrentInstance() on the SeamManagedSessionFactory.

23.8. Spring Application Context as a Seam Component

Although it is possible to use the Spring ContextLoaderListener to start your application's Spring ApplicationContext there are a couple of limitations.

- the Spring ApplicationContext must be started *after* the SeamListener
- it can be tricky starting a Spring ApplicationContext for use in Seam unit and integration tests

To overcome these two limitations the Spring integration includes a Seam component that will start a Spring ApplicationContext. To use this Seam component place the <spring:context-loader/> definition in the components.xml. Specify your Spring context file location in the config-locations attribute. If more than one config file is needed you can place them in the nested <spring:config-locations/> element following standard components.xml multi value practices.

```
<components xmlns="http://jboss.com/products/seam/components"
    xmlns:spring="http://jboss.com/products/seam/spring"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://jboss.com/products/seam/components
        http://jboss.com/products/seam/components-2.0.xsd
        http://jboss.com/products/seam/spring
        http://jboss.com/products/seam/spring-2.0.xsd">

    <spring:context-loader config-locations="/WEB-INF/applicationContext.xml"/>

</components>
```

23.9. Using a Spring TaskExecutor for @Asynchronous

Spring provides an abstraction for executing code asynchronously called a TaskExecutor. The Spring Seam integration allows for the use of a Spring TaskExecutor for executing immediate @Asynchronous method calls. To enable this functionality install the SpringTaskExecutorDispatcher and provide a spring bean defined taskExecutor like so:

```
<spring:task-executor-dispatcher task-executor="#{springThreadPoolTaskExecutor}" />
```

Because a Spring TaskExecutor does not support scheduling of an asynchronous event a fallback Seam Dis-

patcher can be provided to handle scheduled asynchronous event like so:

```
        <!-- Install a ThreadPoolDispatcher to handle scheduled asynchronous event -->
<core:thread-pool-dispatcher name="threadPoolDispatcher"/>

<!-- Install the SpringDispatcher as default -->
<spring:task-executor-dispatcher task-executor="#{springThreadPoolTaskExecutor}" schedule-dispatcher=
```

Chapter 24. Hibernate Search

24.1. Introduction

Full text search engines like Apache Lucene™ are a very powerful technology to bring free text/efficient queries to applications. It suffers several mismatches when dealing with an object domain model (keeping the index up to date, mismatch between the index structure and the domain model, querying mismatch...) Hibernate Search indexes your domain model thanks to a few annotations, takes care of the database / index synchronization and brings you back regular managed objects from free text queries. Hibernate Search is using Apache Lucene under the cover.

Hibernate Search has been designed to integrate nicely and as naturally as possible with JPA and Hibernate. As a natural extension, JBoss Seam provides an Hibernate Search integration.

Please refer to the Hibernate Search documentation [???] for information specific to the Hibernate Search project.

24.2. Configuration

Hibernate Search is configured either in the META-INF/persistence.xml or hibernate.cfg.xml file.

Hibernate Search configuration has sensible defaults for most configuration parameters, Here is a description of the minimal configuration to get started.

```
<persistence-unit name="sample">
  <jta-data-source>java:/DefaultDS</jta-data-source>
  <properties>
    [...]
    <!-- use a file system based index -->
    <property name="hibernate.search.default.directory_provider"
      value="org.hibernate.search.store.FSDirectoryProvider"/>
    <!-- directory where the indexes will be stored -->
    <property name="hibernate.search.default.indexBase"
      value="/Users/prod/apps/dvdstore/dvdindexes"/>
  </properties>
</persistence-unit>
```

If you plan to target Hibernate Annotations or EntityManager 3.2.x (embedded into JBoss AS 4.2.GA), you also need to configure the appropriate event listeners.

```
<persistence-unit name="sample">
  <jta-data-source>java:/DefaultDS</jta-data-source>
  <properties>
    [...]
    <!-- use a file system based index -->
    <property name="hibernate.search.default.directory_provider"
      value="org.hibernate.search.store.FSDirectoryProvider"/>
    <!-- directory where the indexes will be stored -->
    <property name="hibernate.search.default.indexBase"
      value="/Users/prod/apps/dvdstore/dvdindexes"/>

    <property name="hibernate.ejb.event.post-insert"
      value="org.hibernate.search.event.FullTextIndexEventListener"/>
    <property name="hibernate.ejb.event.post-update"
      value="org.hibernate.search.event.FullTextIndexEventListener"/>
    <property name="hibernate.ejb.event.post-delete"
      value="org.hibernate.search.event.FullTextIndexEventListener"/>
  </properties>
</persistence-unit>
```

```
</properties>
</persistence-unit>
```

Note

This step is no longer useful if Hibernate Annotation or EntityManager 3.3.x are used.

In addition to the configuration file, the following jars have to be deployed:

- hibernate-search.jar
- hibernate-commons-annotations.jar
- lucene-core.jar

Note

If you deploy those in a EAR, don't forget to update application.xml

24.3. Usage

Hibernate Search uses annotations to map entities to a Lucene index, check the reference documentation [http://www.hibernate.org/hib_docs/search/reference/en/html_single/] for more informations.

Hibernate Search is fully integrated with the API and semantic of JPA / Hibernate. Switching from a HQL or Criteria based query requires just a few lines of code. The main API the application interacts with is the Full-TextSession API (subclass of Hibernate's Session).

When Hibernate Search is present, JBoss Seam injects a FullTextSession.

```
@Stateful
@Name("search")
public class FullTextSearchAction implements FullTextSearch, Serializable
{
    @In
    FullTextSession session;

    public void search(String searchString) {
        org.apache.lucene.query.Query luceneQuery = getLuceneQuery();
        org.hibernate.Query query session.createFullTextQuery(luceneQuery, Product.class);
        searchResults = query
            .setMaxResults(pageSize + 1)
            .setFirstResult(pageSize * currentPage)
            .list();
    }
    [...]
}
```

Note

FullTextSession extends org.hibernate.Session so that it can be used as a regular Hibernate Session

If the Java Persistence API is used, a smoother integration is proposed.

```
@Name("search")
public class FullTextSearchAction implements FullTextSearch, Serializable
```

```

{
    @In
    FullTextEntityManager em;

    public void search(String searchString) {
        org.apache.lucene.query.Query luceneQuery = getLuceneQuery();
        javax.persistence.Query query = em.createFullTextQuery(luceneQuery, Product.class);
        searchResults = query
            .setMaxResults(pageSize + 1)
            .setFirstResult(pageSize * currentPage)
            .getResultList();
    }
    [...]
}
    
```

When Hibernate Search is present, a `FullTextEntityManager` is injected. `FullTextEntityManager` extends `EntityManager` with search specific methods, the same way `FullTextSession` extends `Session`.

When an EJB 3.0 Session or Message Driven Bean injection is used (ie Bean using `@PersistenceContext`), it is not possible to replace the `EntityManager` interface by the `FullTextEntityManager` interface in the declaration statement. However, the implementation injected will be a `FullTextEntityManager` implementation: down-casting is then possible.

```

@Stateful
@Name("search")
public class FullTextSearchAction implements FullTextSearch, Serializable
{
    @PersistenceContext
    EntityManager em;

    public void search(String searchString) {
        org.apache.lucene.query.Query luceneQuery = getLuceneQuery();
        FullTextEntityManager ftEm = (FullTextEntityManager) em;
        javax.persistence.Query query = ftEm.createFullTextQuery(luceneQuery, Product.class);
        searchResults = query
            .setMaxResults(pageSize + 1)
            .setFirstResult(pageSize * currentPage)
            .getResultList();
    }
    [...]
}
    
```

Caution

For people accustomed to Hibernate Search out of Seam, note that using `Search.createFullTextSession` is not necessary.

Check the `DVDStore` or the blog examples of the JBoss Seam distribution for a concrete use of Hibernate Search.

Chapter 25. Configuring Seam and packaging Seam applications

Configuration is a very boring topic and an extremely tedious pastime. Unfortunately, several lines of XML are required to integrate Seam into your JSF implementation and servlet container. There's no need to be too put off by the following sections; you'll never need to type any of this stuff yourself, since you can just copy and paste from the example applications!

25.1. Basic Seam configuration

First, let's look at the basic configuration that is needed whenever we use Seam with JSF.

25.1.1. Integrating Seam with JSF and your servlet container

Of course, you need a faces servlet!

```
<servlet>
  <servlet-name>Faces Servlet</servlet-name>
  <servlet-class>javax.faces.webapp.FacesServlet</servlet-class>
  <load-on-startup>1</load-on-startup>
</servlet>

<servlet-mapping>
  <servlet-name>Faces Servlet</servlet-name>
  <url-pattern>*.seam</url-pattern>
</servlet-mapping>
```

(You can adjust the URL pattern to suit your taste.)

In addition, Seam requires the following entry in your `web.xml` file:

```
<listener>
  <listener-class>org.jboss.seam.servlet.SeamListener</listener-class>
</listener>
```

This listener is responsible for bootstrapping Seam, and for destroying session and application contexts.

Some JSF implementations have a broken implementation of server-side state saving that interferes with Seam's conversation propagation. If you have problems with conversation propagation during form submissions, try switching to client-side state saving. You'll need this in `web.xml`:

```
<context-param>
  <param-name>javax.faces.STATE_SAVING_METHOD</param-name>
  <param-value>client</param-value>
</context-param>
```

25.1.2. Using facelets

If you want follow our advice and use facelets instead of JSP, add the following lines to `faces-config.xml`:

```
<application>
  <view-handler>com.sun.facelets.FaceletViewHandler</view-handler>
</application>
```

And the following lines to `web.xml`:

```
<context-param>
  <param-name>javax.faces.DEFAULT_SUFFIX</param-name>
  <param-value>.xhtml</param-value>
</context-param>
```

25.1.3. Seam Resource Servlet

The Seam Resource Servlet provides resources used by Seam Remoting, captchas (see the security chapter) and some JSF UI controls. Configuring the Seam Resource Servlet requires the following entry in `web.xml`:

```
<servlet>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <servlet-class>org.jboss.seam.servlet.SeamResourceServlet</servlet-class>
</servlet>

<servlet-mapping>
  <servlet-name>Seam Resource Servlet</servlet-name>
  <url-pattern>/seam/resource/*</url-pattern>
</servlet-mapping>
```

25.1.4. Seam servlet filters

Seam doesn't need any servlet filters for basic operation. However, there are several features which depend upon the use of filters. To make things easier, Seam lets you add and configure servlet filters just like you would configure other built-in Seam components. To take advantage of this feature, we must first install a master filter in `web.xml`:

```
<filter>
  <filter-name>Seam Filter</filter-name>
  <filter-class>org.jboss.seam.servlet.SeamFilter</filter-class>
</filter>

<filter-mapping>
  <filter-name>Seam Filter</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
```

The Seam master filter *must* be the first filter specified in `web.xml`. This ensures it is run first.

To disable a built in filter, you can set `disabled="true"` on a particular filter.

Adding the master filter enables the following built-in filters.

Exception handling

This filter provides the exception mapping functionality in `pages.xml` (almost all applications will need this). It also takes care of rolling back uncommitted transactions when uncaught exceptions occur. (According to the Java EE specification, the web container should do this automatically, but we've found that this behavior cannot be relied upon in all application servers. And it is certainly not required of plain servlet engines like Tomcat.)

By default, the exception handling filter will process all requests, however this behavior may be adjusted by adding a `<web:exception-filter>` entry to `components.xml`, as shown in this example:

```
<components xmlns="http://jboss.com/products/seam/components"
  xmlns:web="http://jboss.com/products/seam/web">
```

```
<web:exception-filter url-pattern="*.seam"/>
</components>
```

- `url-pattern` — Used to specify which requests are filtered, the default is all requests.

Conversation propagation with redirects

This filter allows Seam to propagate the conversation context across browser redirects. It intercepts any browser redirects and adds a request parameter that specifies the Seam conversation identifier.

The redirect filter will process all requests by default, but this behavior can also be adjusted in `components.xml`:

```
<web:redirect-filter url-pattern="*.seam"/>
```

- `url-pattern` — Used to specify which requests are filtered, the default is all requests.

Multipart form submissions

This feature is necessary when using the Seam file upload JSF control. It detects multipart form requests and processes them according to the multipart/form-data specification (RFC-2388). To override the default settings, add the following entry to `components.xml`:

```
<web:multipart-filter create-temp-files="true"
                      max-request-size="1000000"
                      url-pattern="*.seam"/>
```

- `create-temp-files` — If set to `true`, uploaded files are written to a temporary file (instead of held in memory). This may be an important consideration if large file uploads are expected. The default setting is `false`.
- `max-request-size` — If the size of a file upload request (determined by reading the `Content-Length` header in the request) exceeds this value, the request will be aborted. The default setting is 0 (no size limit).
- `url-pattern` — Used to specify which requests are filtered, the default is all requests.

Character encoding

Sets the character encoding of submitted form data.

This filter is not installed by default and requires an entry in `components.xml` to enable it:

```
<web:character-encoding-filter encoding="UTF-16"
                              override-client="true"
                              url-pattern="*.seam"/>
```

- `encoding` — The encoding to use.
- `override-client` — If this is set to `true`, the request encoding will be set to whatever is specified by `encoding` no matter whether the request already specifies an encoding or not. If set to `false`, the request encoding will only be set if the request doesn't already specify an encoding. The default setting is `false`.

- `url-pattern` — Used to specify which requests are filtered, the default is all requests.

RichFaces

If RichFaces is used in your project, Seam will install the RichFaces Ajax filter for you, making sure to install it before all other built-in filters. You don't need to install the RichFaces Ajax filter in `web.xml` yourself.

The RichFaces Ajax filter is only installed if the RichFaces jars are present in your project.

To override the default settings, add the following entry to `components.xml`. The options are the same as those specified in the RichFaces Developer Guide:

```
<web:ajax4jsf-filter force-parser="true"
                    enable-cache="true"
                    log4j-init-file="custom-log4j.xml"
                    url-pattern="*.seam"/>
```

- `force-parser` — forces all JSF pages to be validated by Richfaces's XML syntax checker. If `false`, only AJAX responses are validated and converted to well-formed XML. Setting `force-parser` to `false` improves performance, but can provide visual artifacts on AJAX updates.
- `enable-cache` — enables caching of framework-generated resources (e.g. javascript, CSS, images, etc). When developing custom javascript or CSS, setting to `true` prevents the browser from caching the resource.
- `log4j-init-file` — is used to setup per-application logging. A path, relative to web application context, to the `log4j.xml` configuration file should be provided.
- `url-pattern` — Used to specify which requests are filtered, the default is all requests.

Identity Logging

This filter adds the authenticated user name to the `log4j` mapped diagnostic context so that it can be included in formatted log output if desired, by adding `%X{username}` to the pattern.

By default, the logging filter will process all requests, however this behavior may be adjusted by adding a `<web:logging-filter>` entry to `components.xml`, as shown in this example:

```
<components xmlns="http://jboss.com/products/seam/components"
            xmlns:web="http://jboss.com/products/seam/web">
  <web:logging-filter url-pattern="*.seam"/>
</components>
```

- `url-pattern` — Used to specify which requests the filter is active for. The default is all requests.

Context management for custom servlets

Requests sent direct to some servlet other than the JSF servlet are not processed through the JSF lifecycle, so Seam provides a servlet filter that can be applied to any other servlet that needs access to Seam components.

This filter allows custom servlets to interact with the Seam contexts. It sets up the Seam contexts at the beginning of each request, and tears them down at the end of the request. You should make sure that this filter is *never* applied to the JSF `FacesServlet`. Seam uses the phase listener for context management in a JSF request.

This filter is not installed by default and requires an entry in `components.xml` to enable it:

```
<web:context-filter url-pattern="/media/*"/>
```

- `url-pattern` — Used to specify which requests are filtered, the default is all requests. If the `url-pattern` is specified for the context filter, then the filter will be enabled (unless explicitly disabled).

The context filter expects to find the conversation id of any conversation context in a request parameter named `conversationId`. You are responsible for ensuring that it gets sent in the request.

You are also responsible for ensuring propagation of any new conversation id back to the client. Seam exposes the conversation id as a property of the built in component conversation.

Adding custom filters

Seam can install your filters for you, allowing you to specify *where* in the chain your filter is placed (the servlet specification doesn't provide a well defined order if you specify your filters in a `web.xml`). Just add the `@Filter` annotation to your Seam component (which must implement `javax.servlet.Filter`):

```
@Startup
@Scope(APPLICATION)
@Name("org.jboss.seam.web.multipartFilter")
@BypassInterceptors
@Filter(within="org.jboss.seam.web.ajax4jsfFilter")
public class MultipartFilter extends AbstractFilter {
```

Adding the `@Startup` annotation means that the component is available during Seam startup; bijection isn't available here (`@BypassInterceptors`); and the filter should be further down the chain than the RichFaces filter (`@Filter(within="org.jboss.seam.web.ajax4jsfFilter")`).

25.1.5. Integrating Seam with your EJB container

We need to apply the `SeamInterceptor` to our Seam components. The simplest way to do this across an entire application is to add the following interceptor configuration in `ejb-jar.xml`:

```
<interceptors>
  <interceptor>
    <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
  </interceptor>
</interceptors>

<assembly-descriptor>
  <interceptor-binding>
    <ejb-name>*</ejb-name>
    <interceptor-class>org.jboss.seam.ejb.SeamInterceptor</interceptor-class>
  </interceptor-binding>
</assembly-descriptor>
```

Seam needs to know where to go to find session beans in JNDI. One way to do this is specify the `@JndiName` annotation on every session bean Seam component. However, this is quite tedious. A better approach is to specify a pattern that Seam can use to calculate the JNDI name from the EJB name. Unfortunately, there is no standard mapping to global JNDI defined in the EJB3 specification, so this mapping is vendor-specific. We usually specify this option in `components.xml`.

For JBoss AS, the following pattern is correct:

```
<core:init jndi-name="myEarName/#{ejbName}/local" />
```

Where `myEarName` is the name of the EAR in which the bean is deployed.

Outside the context of an EAR (when using the JBoss Embeddable EJB3 container), the following pattern is the one to use:

```
<core:init jndi-name="#{ejbName}/local" />
```

You'll have to experiment to find the right setting for other application servers. Note that some servers (such as GlassFish) require you to specify JNDI names for all EJB components explicitly (and tediously). In this case, you can pick your own pattern ;-)

In an EJB3 environment, we recommend the use of a special built-in component for transaction management, that is fully aware of container transactions, and can correctly process transaction success events registered with the Events component. If you don't add this line to your `components.xml` file, Seam won't know when container-managed transactions end:

```
<transaction:ejb-transaction/>
```

25.1.6. Don't forget!


There is one final item you need to know about. You must place a `seam.properties`, `META-INF/seam.properties` or `META-INF/components.xml` file in any archive in which your Seam components are deployed (even an empty properties file will do). At startup, Seam will scan any archives with `seam.properties` files for seam components.

In a web archive (WAR) file, you must place a `seam.properties` file in the `WEB-INF/classes` directory if you have any Seam components included here.

That's why all the Seam examples have an empty `seam.properties` file. You can't just delete this file and expect everything to still work!

You might think this is silly and what kind of idiot framework designers would make an empty file affect the behavior of their software?? Well, this is a workaround for a limitation of the JVM—if we didn't use this mechanism, our next best option would be to force you to list every component explicitly in `components.xml`, just like some other competing frameworks do! I think you'll like our way better.

25.2. Configuring Seam in Java EE 5



If you're running in a Java EE 5 environment, this is all the configuration required to start using Seam!

25.2.1. Packaging

Once you've packaged all this stuff together into an EAR, the archive structure will look something like this:

```
my-application.ear/  
  jboss-seam.jar  
  lib/  
    jboss-el.jar  
  META-INF/  
    MANIFEST.MF  
    application.xml  
  my-application.war/  
    META-INF/  
      MANIFEST.MF  
    WEB-INF/  
      web.xml  
      components.xml  
      faces-config.xml  
      lib/  
        jsf-facelets.jar  
        jboss-seam-ui.jar  
      login.jsp  
      register.jsp  
      ...  
  my-application.jar/  
    META-INF/  
      MANIFEST.MF  
      persistence.xml  
    seam.properties  
  org/  
    jboss/  
      myapplication/  
        User.class  
        Login.class  
        LoginBean.class  
        Register.class  
        RegisterBean.class  
        ...
```

You should declare `jboss-seam.jar` as an `ejb` module in `META-INF/application.xml`; `jboss-el.jar` should be placed in the EAR's `lib` directory (putting it in the EAR classpath).

If you want to use `jBPM` or `Drools`, you must include the needed jars in the EAR's `lib` directory.

If you want to use `facelets` (our recommendation), you must include `jsf-facelets.jar` in the `WEB-INF/lib` directory of the WAR.

If you want to use the Seam tag library (most Seam applications do), you must include `jboss-seam-ui.jar` in the `WEB-INF/lib` directory of the WAR. If you want to use the PDF or email tag libraries, you need to put `jboss-seam-pdf.jar` or `jboss-seam-mail.jar` in `WEB-INF/lib`.


If you want to use the Seam debug page (only works for applications using `facelets`), you must include `jboss-seam-debug.jar` in the `WEB-INF/lib` directory of the WAR.

Seam ships with several example applications that are deployable in any Java EE container that supports EJB 3.0.

I really wish that was all there was to say on the topic of configuration but unfortunately we're only about a third of the way there. If you're too overwhelmed by all this tedious configuration stuff, feel free to skip over the rest of this section and come back to it later.

25.3. Configuring Seam in J2EE

Seam is useful even if you're not yet ready to take the plunge into EJB 3.0. In this case you would use Hibernate3 or JPA instead of EJB 3.0 persistence, and plain JavaBeans instead of session beans. You'll miss out on some of the nice features of session beans but it will be very easy to migrate to EJB 3.0 when you're ready and, in the meantime, you'll be able to take advantage of Seam's unique declarative state management architecture.



Seam JavaBean components do not provide declarative transaction demarcation like session beans do. You *could* manage your transactions manually using the JTA `UserTransaction` or declaratively using Seam's `@Transactional` annotation. But most applications will just use Seam managed transactions when using Hibernate with JavaBeans.

The Seam distribution includes a version of the booking example application that uses Hibernate3 and JavaBeans instead of EJB3, and another version that uses JPA and JavaBeans. These example applications are ready to deploy into any J2EE application server.

25.3.1. Bootstrapping Hibernate in Seam

Seam will bootstrap a Hibernate `SessionFactory` from your `hibernate.cfg.xml` file if you install a built-in component:

```
<persistence:hibernate-session-factory name="hibernateSessionFactory"/>
```

You will also need to configure a *managed session* if you want a Seam managed Hibernate Session to be available via injection.

```
<persistence:managed-hibernate-session name="hibernateSessionFactory"
    session-factory="#{hibernateSessionFactory}"/>
```

25.3.2. Bootstrapping JPA in Seam

Seam will bootstrap a JPA `EntityManagerFactory` from your `persistence.xml` file if you install this built-in component:

```
<persistence:entity-manager-factory name="entityManagerFactory"/>
```

You will also need to configure a *managed persistence context* if you want a Seam managed JPA EntityManager to be available via injection.

```
<persistence:managed-persistence-context name="entityManager"
    entity-manager-factory="#{entityManagerFactory}"/>
```

25.3.3. Packaging

We can package our application as a WAR, in the following structure:

```
my-application.war/
  META-INF/
    MANIFEST.MF
  WEB-INF/
    web.xml
    components.xml
    faces-config.xml
    lib/
      jboss-seam.jar
      jboss-seam-ui.jar
      jboss-el.jar
      jsf-facelets.jar
      hibernate3.jar
      hibernate-annotations.jar
      hibernate-validator.jar
      ...
    my-application.jar/
      META-INF/
        MANIFEST.MF
      seam.properties
      hibernate.cfg.xml
      org/
        jboss/
          myapplication/
            User.class
            Login.class
            Register.class
            ...
    login.jsp
    register.jsp
    ...
```

If we want to deploy Hibernate in a non-EE environment like Tomcat or TestNG, we need to do a little bit more work.

25.4. Configuring Seam in Java SE, without JBoss Embedded

It is possible to use Seam completely outside of an EE environment. In this case, you need to tell Seam how to manage transactions, since there will be no JTA available. If you're using JPA, you can tell Seam to use JPA resource-local transactions, ie. EntityTransaction, like so:

```
<transaction:entity-transaction entity-manager="#{entityManager}"/>
```

If you're using Hibernate, you can tell Seam to use the Hibernate transaction API like this:

```
<transaction:hibernate-transaction session="#{session}"/>
```


Of course, you'll also need to define a datasource.

A better alternative is to use JBoss Embedded to get access to the EE APIs.


25.5. Configuring Seam in Java SE, with JBoss Embedded

JBoss Embedded lets you run EJB3 components outside the context of the Java EE 5 application server. This is especially, but not only, useful for testing.

The Seam booking example application includes a TestNG integration test suite that runs on JBoss Embedded via SeamTest.



The booking example application may even be deployed to Tomcat.



25.5.1. Installing Embedded JBoss

Embedded JBoss must be installed into Tomcat for Seam applications to run correctly on it. Embedded JBoss only runs on JDK 1.5 (not JDK 1.6). Embedded JBoss can be downloaded here [http://sourceforge.net/project/showfiles.php?group_id=22866&package_id=228977]. The process for installing Embedded JBoss into Tomcat 6 is quite simple. First, you should copy the Embedded JBoss JARs and configuration files into Tomcat.

- Copy all files and directories under the Embedded JBoss bootstrap and lib directories, except for the `jndi.properties` file, into the Tomcat lib directory.
- Remove the `annotations-api.jar` file from the Tomcat lib directory.

Next, two configuration files need to be updated to add Embedded JBoss-specific functionality.

- Add the Embedded JBoss listener to `conf/server.xml`. It should appear after all other listeners in the file.

```
<Listener className="org.jboss.embedded.tomcat.EmbeddedJBossBootstrapListener" />
```

- WAR file scanning should be enabled by adding a listener to `conf/context.xml`.

```
<Listener className="org.jboss.embedded.tomcat.WebinfScanner" />
```

For more configuration options, please see the Embedded JBoss Tomcat integration wiki entry [<http://wiki.jboss.org/wiki/Wiki.jsp?page=EmbeddedAndTomcat>].

25.5.2. Packaging

The archive structure of a WAR-based deployment on an servlet engine like Tomcat will look something like this:

```
my-application.war/
  META-INF/
    MANIFEST.MF
  WEB-INF/
    web.xml
    components.xml
    faces-config.xml
    lib/
      jboss-seam.jar
      jboss-seam-ui.jar
      jboss-el.jar
      jsf-facelets.jar
      jsf-api.jar
      jsf-impl.jar
      ...
    my-application.jar/
      META-INF/
        MANIFEST.MF
        persistence.xml
      seam.properties
      org/
        jboss/
          myapplication/
            User.class
            Login.class
            LoginBean.class
            Register.class
            RegisterBean.class
            ...
    login.jsp
    register.jsp
    ...
```

Most of the Seam example applications may be deployed to Tomcat by running `ant deploy.tomcat`.

25.6. Configuring jBPM in Seam

Seam's jBPM integration is not installed by default, so you'll need to enable jBPM by installing a built-in component. You'll also need to explicitly list your process and pageflow definitions. In `components.xml`:

```
<bpm:jbp>
  <bpm:pageflow-definitions>
    <value>createDocument.jpdl.xml</value>
    <value>editDocument.jpdl.xml</value>
    <value>approveDocument.jpdl.xml</value>
```

```

    </bpm:pageflow-definitions>
    <bpm:process-definitions>
      <value>documentLifecycle.jpdl.xml</value>
    </bpm:process-definitions>
  </bpm:jbpm>

```

No further special configuration is needed if you only have pageflows. If you do have business process definitions, you need to provide a jBPM configuration, and a Hibernate configuration for jBPM. The Seam DVD Store demo includes example `jbpm.cfg.xml` and `hibernate.cfg.xml` files that will work with Seam:

```

<jbpm-configuration>
  <jbpm-context>
    <service name="persistence">
      <factory>
        <bean class="org.jbpm.persistence.db.DbPersistenceServiceFactory">
          <field name="isTransactionEnabled"><false/></field>
        </bean>
      </factory>
    </service>
    <service name="tx" factory="org.jbpm.tx.TxServiceFactory" />
    <service name="message" factory="org.jbpm.msg.db.DbMessageServiceFactory" />
    <service name="scheduler" factory="org.jbpm.scheduler.db.DbSchedulerServiceFactory" />
    <service name="logging" factory="org.jbpm.logging.db.DbLoggingServiceFactory" />
    <service name="authentication"
      factory="org.jbpm.security.authentication.DefaultAuthenticationServiceFactory" />
  </jbpm-context>
</jbpm-configuration>

```

The most important thing to notice here is that jBPM transaction control is disabled. Seam or EJB3 should control the JTA transactions.

25.6.1. Packaging

There is not yet any well-defined packaging format for jBPM configuration and process/pageflow definition files. In the Seam examples we've decided to simply package all these files into the root of the EAR. In future, we will probably design some other standard packaging format. So the EAR looks something like this:

```

my-application.ear/
  jboss-seam.jar
  lib/
    jboss-el.jar
    jbpm-3.1.jar
  META-INF/
    MANIFEST.MF
    application.xml
  my-application.war/
    META-INF/
      MANIFEST.MF
    WEB-INF/
      web.xml
      components.xml
      faces-config.xml
      lib/
        jsf-facelets.jar
        jboss-seam-ui.jar
      login.jsp
      register.jsp
      ...
  my-application.jar/
    META-INF/
      MANIFEST.MF
      persistence.xml
    seam.properties

```

```

org/
  jboss/
    myapplication/
      User.class
      Login.class
      LoginBean.class
      Register.class
      RegisterBean.class
      ...
jbpm.cfg.xml
hibernate.cfg.xml
createDocument.jpdl.xml
editDocument.jpdl.xml
approveDocument.jpdl.xml
documentLifecycle.jpdl.xml

```

25.7. Configuring Seam in a Portal

To run a Seam application as a portlet, you'll need to provide certain portlet metadata (`portlet.xml`, etc) in addition to the usual Java EE metadata. See the `examples/portal` directory for an example of the booking demo preconfigured to run on JBoss Portal.

25.8. Configuring SFSB and Session Timeouts in JBoss AS

It is very important that the timeout for Stateful Session Beans is set higher than the timeout for HTTP Sessions, otherwise SFSB's may time out before the user's HTTP session has ended. JBoss Application Server has a default session bean timeout of 30 minutes, which is configured in `server/default/conf/standardjboss.xml` (replace *default* with your own configuration).

The default SFSB timeout can be adjusted by modifying the value of `max-bean-life` in the `LRUStatefulContextCachePolicy` cache configuration:

```

<container-cache-conf>
  <cache-policy>org.jboss.ejb.plugins.LRUStatefulContextCachePolicy</cache-policy>
  <cache-policy-conf>
    <min-capacity>50</min-capacity>
    <max-capacity>1000000</max-capacity>
    <remover-period>1800</remover-period>

    <!-- SFSB timeout in seconds; 1800 seconds == 30 minutes -->
    <max-bean-life>1800</max-bean-life>

    <overager-period>300</overager-period>
    <max-bean-age>600</max-bean-age>
    <resizer-period>400</resizer-period>
    <max-cache-miss-period>60</max-cache-miss-period>
    <min-cache-miss-period>1</min-cache-miss-period>
    <cache-load-factor>0.75</cache-load-factor>
  </cache-policy-conf>
</container-cache-conf>

```

The default HTTP session timeout can be modified in `server/default/deploy/jbossweb-tomcat55.sar/conf/web.xml` for JBoss 4.0.x, or in `server/default/deploy/jboss-web.deployer/conf/web.xml` for JBoss 4.2.x. The following entry in this file controls the default session timeout for all web applications:

```

<session-config>
  <!-- HTTP Session timeout, in minutes -->
  <session-timeout>30</session-timeout>

```

```
</session-config>
```

To override this value for your own application, simply include this entry in your application's own `web.xml`.

Chapter 26. Seam on OC4J

OC4J (Oracle Containers for Java) 11g (currently a "Technology Preview" release) is Oracle's JEE 5 application server. We will start by looking at the building and deploying the Hotel Booking example application which comes with Seam, and then at deploying a project generated by seam-gen. This project will integrate Seam, RichFaces Ajax and components, Seam Security (with Drools), Facelets and JPA provided by Hibernate.

This section requires you to use OC4J 11g Technology Preview (not OC4J 10g). You can download OC4J 11g from <http://www.oracle.com/technology/tech/java/oc4j/11/>.

26.1. The jee5/booking example

The jee5/booking example is based on the Hotel Booking example (which runs on JBoss AS). Out of the box it is designed to run on Glassfish, but it's easy to build it for OC4J.

26.1.1. Booking Example Dependencies

First, let's look at the dependencies of the booking example. Armed with this knowledge we can look at the extra dependencies requirements that OC4J adds.

- `jboss-seam.jar` — We declare this as an EJB3 module (why? well Seam needs to be able to interact with container managed transactions; this is implemented as an EJB3 Stateful Session Bean)
- `jboss-el.jar`
- `jboss-seam-ui.jar` — Seam's JSF controls depend on Apache's commons-beanutils
- `jboss-seam-debug.jar`
- `jsf-facelets.jar`
- `richfaces-api.jar` — which requires Apache commons-digester and commons-beanutils
- `richfaces-impl.jar` and `richfaces-ui.jar` — which requires Apache commons-digester and commons-beanutils

26.1.2. Extra dependencies required by OC4J

- Hibernate — of course, we decided to use Hibernate as the JPA provider (rather than TopLink Essentials which ships with OC4J).

To use Hibernate as your JPA provider you need three jars (`hibernate3.jar`, `hibernate-annotations.jar`, `hibernate-entitymanager.jar`) and their dependencies (`jboss-common.jar`, `jboss-archive-browsing.jar` and `ejb3-persistence.jar`). You can find these in the `hibernate/lib` directory in the Seam distribution.

- `thirdparty-all.jar` — a selection of third party libraries on which Seam depends (like `javassist`).

Running Seam on most application servers (such as JBoss AS or Glassfish) you only need to include the dependencies for those bits of Seam you actually use (e.g. if you use Seam Text you need to include ANTLR);

but, on OC4J, due to its "interesting" classloading you must always include them:

- antlr-2.7.6.jar — needed for Seam Text (not used in the example).
- jbpmm-jpdl.jar — needed for Seam's JBPM integration (not used in the example).
- Drools — needed for Seam Security. We aren't using Seam security with Drools, but have to include it. Drools consists of 5 jars - drools-core-4.0.0.jar, drools-compiler-4.0.0.jar, janino-2.5.7.jar, mvel14-1.2rc1.jar and antlr-runtime-3.0.jar. Drools integration is not used in the example.

26.1.3. Configuration file changes

There are just a few changes to be made:

web.xml

you need to declare all your ejb's in the web.xml. This is a silly requirement of a number of JEE 5 application servers - for example OC4J and Glassfish.

```
<ejb-local-ref>
  <ejb-ref-name>
    jboss-seam-jee5/AuthenticatorAction/local
  </ejb-ref-name>
  <ejb-ref-type>Session</ejb-ref-type>
  <local-home/>
  <local>
    org.jboss.seam.example.booking.Authenticator
  </local>
  <ejb-link>AuthenticatorAction</ejb-link>
</ejb-local-ref>
```

persistence.xml

you need to provide the correct configuration for your JPA implementation. We are using Hibernate and due to OC4J bundling an old ANTLR, we need to use an alternative query factory, we also want to use the OC4J transaction manager:

```
<property
  name="hibernate.query.factory_class"
  value="org.hibernate.hql.classic.ClassicQueryTranslatorFactory" />
<property
  name="hibernate.transaction.manager_lookup_class"
  value="org.hibernate.transaction.OrionTransactionManagerLookup" />
```

26.1.4. Building the jee5/booking example

1. Modify the following files in the project:

- build.xml — Un-comment the OC4J-related libraries
- resources/META-INF/persistence.xml — Comment out the Glassfish properties and un-comment the OC4J properties.

2. Build the demo app by running ant. The build target is dist/jboss-seam-jee5.ear

3. Copy hsqldb.jar to OC4J: cp ../../seam-gen/lib/hsqldb.jar \$ORACLE_HOME/j2ee/home/applib/

(OC4J doesn't come with an embedded database so we decided to use HSQLDB)

26.2. Deploying a Seam application to OC4J

This mini-tutorial describes the (fairly tedious) steps required to deploy a JEE 5 application to OC4J. It assumes you are deploying the `jee5/booking` example, using the embedded `hsqldb` database. To deploy another application you would need to alter the `datasource` and application name.

1. Download and unzip OC4J
2. Make sure you have `$JAVA_HOME` and `$ORACLE_HOME` set as environment variables (`$ORACLE_HOME` is the directory to which you unzip OC4J). For further information on installing OC4J, consult the `Readme.txt` distributed with OC4J
3. Edit the OC4J `datasource` `$ORACLE_HOME/j2ee/home/config/data-sources.xml` and, inside `<data-sources>`, add

```
<managed-data-source
  connection-pool-name="jee5-connection-pool"
  jndi-name="jdbc/__default"
  name="jee5-managed-data-source" />
<connection-pool name="jee5-connection-pool">
  <connection-factory
    factory-class="org.hsqldb.jdbcDriver"
    user="sa"
    password=""
    url="jdbc:hsqldb:." />
</connection-pool>
```

The `jndi-name` is used as the `jta-data-source` in `persistence.xml`.

4. Edit `$ORACLE_HOME/j2ee/home/config/server.xml` and, inside `<application-server>`, add

```
<application name="jboss-seam-jee5"
  path="../../../home/applications/jboss-seam-jee5.ear"
  parent="default"
  start="true" />
```

To keep things simple use the same names as you used for project.

5. Edit `$ORACLE_HOME/j2ee/home/config/default-web-site.xml`, and, inside `<web-site>`, add

```
<web-app application="jboss-seam-jee5"
  name="jboss-seam-jee5"
  load-on-startup="true"
  root="/seam-jee5" />
```

The root is the context path you will put into your web browser to access the application.

6. Copy the application to OC4J: `cp dist/jboss-seam-jee5.ear $ORACLE_HOME/j2ee/home/applications/`
7. Start OC4J: `$ORACLE_HOME/j2ee/home/java -jar -XX:MaxPermSize=256M oc4j.jar`

You must override the default PermGen memory settings using above command. See <http://www.oracle.com/technology/tech/java/oc4j/11/oc4j-relnotes.html> for details.

You will be asked to set the admin password if this is the first time you've started OC4J

8. Checkout the app at: <http://localhost:8888/seam-jee5>
9. You can stop the server by pressing CTRL-C in the console on which the server is running.

26.3. Deploying an application created using seam-gen to OC4J

The following explanation assumes you are using the command line and a simple text editor, but of course you can use your favourite IDE - seam-gen projects come with support for Eclipse and Netbeans.

We start by creating a pretty simple application using seam-gen. seam-gen uses Hibernate Tools to reverse engineer a database schema to JPA entity beans; it also generates Seam Application Framework components and JSF views for CRUD. This tutorial uses MySQL (but of course you could use any database, altering the SQL as appropriate); install, configure and run MySQL, then create a database with some sample data.

Next, run `./seam setup` in the seam directory.

```
> ./seam setup
Buildfile: build.xml

setup:
[echo] Welcome to seam-gen :-)
[input] Enter your Java project workspace (the directory that contains your Seam projects) [/home/
[input] Enter your JBoss home directory [/home/pmuir/java/jboss-4.2.1.GA] [/home/pmuir/java/jboss-
[input] Enter the project name [oc4j-example] [oc4j-example]
[input] Is this project deployed as an EAR (with EJB components) or a WAR (with no EJB support) [
[input] Enter the Java package name for your session beans [org.jboss.seam.tutorial.oc4j.action] [
[input] Enter the Java package name for your entity beans [org.jboss.seam.tutorial.oc4j.model] [or
[input] Enter the Java package name for your test cases [org.jboss.seam.tutorial.oc4j.test] [org.
[input] What kind of database are you using? [mysql] (hsqldb, [mysql], oracle, postgres, mssql, db2)
[input] Enter the Hibernate dialect for your database [org.hibernate.dialect.MySQLDialect] [org.h
[input] Enter the filesystem path to the JDBC driver jar [lib/mysql.jar] [lib/mysql.jar]
[input] Enter JDBC driver class for your database [com.mysql.jdbc.Driver] [com.mysql.jdbc.Driver]
[input] Enter the JDBC URL for your database [jdbc:mysql:///oc4j] [jdbc:mysql:///oc4j]
[input] Enter database username [user] [user]
[input] Enter database password [password] [password]
[input] skipping input as property hibernate.default_schema.new has already been set.
[input] Enter the database catalog name (it is OK to leave this blank) [] []
[input] Are you working with tables that already exist in the database? [y] ([y], n, )
[input] Do you want to drop and recreate the database tables and data in import.sql each time you
[propertyfile] Updating property file: /home/pmuir/workspace/jboss-seam/seam-gen/build.properties
[echo] Installing JDBC driver jar to JBoss server
[echo] Type 'seam new-project' to create the new project

BUILD SUCCESSFUL
```

Type `./seam new-project` to create your project and `cd` to the newly created project.

Type `./seam generate-entities` to run create the entities, the Seam Application Framework classes and the relevant views.

We now need to make some changes to the generated project. Let's start with the configuration files:

`resources/META-INF/persistence-dev.xml`

- Alter the `jta-data-source` to be `jdbc/__oc4jExample` (and use this as the `jndi-name` when creating the data source in `data-sources.xml`).
- Add the properties (described above):

```
<property name="hibernate.query.factory_class"
  value="org.hibernate.hql.classic.ClassicQueryTranslatorFactory" />
<property name="hibernate.transaction.manager_lookup_class"
  value="org.hibernate.transaction.OrionTransactionManagerLookup" />
<property name="hibernate.transaction.flush_before_completion"
  value="true"/>
<property name="hibernate.cache.provider_class"
  value="org.hibernate.cache.HashtableCacheProvider"/>
```

- Remove the JBoss AS specific method of exposing the `EntityManagerFactory`:

```
<property
  name="jboss.entity.manager.factory.jndi.name"
  value="java:/oc4j-exampleEntityManagerFactory">
```

- You'll need to alter `persistence-prod.xml` as well if you want to deploy to OC4J using the prod profile.

`resources/META-INF/jboss-app.xml`

You can delete this file as we aren't deploying to JBoss AS (`jboss-app.xml` is used to enable classloading isolation in JBoss AS)

`resources/*-ds.xml`

You can delete these file as we aren't deploying to JBoss AS (these files define datasources in JBoss AS, in OC4J you have to edit the master `data-sources.xml` file)

`resources/WEB-INF/components.xml`

- Enable container managed transaction integration - add the `<transaction:ejb-transaction />` component, and it's namespace declaration `xmlns:transaction="http://jboss.com/products/seam/transaction"`
- Alter the `jndi-pattern` to `java:comp/env/oc4j-example/#{ejbName}/local`
- We want to use a Seam Managed Persistence Context in our application. Unfortunately OC4J doesn't expose the `EntityManagerFactory` in JNDI, but Seam provides a built-in manager component:

```
<persistence:entity-manager-factory
  auto-create="true"
  name="oc4jEntityManagerFactory"
  persistence-unit-name="oc4j-example" />
```

We then need to tell Seam to use it, so we alter the managed-persistence-context injecting the Entity Manager Factory:

```
<persistence:managed-persistence-context
name="entityManager"
auto-create="true"
entity-manager-factory="{oc4jEntityManagerFactory}" />
```

resources/WEB-INF/web.xml

You need to declare all your EJBs here. Remember to include the Seam container managed transaction integration:

```
<ejb-local-ref>
  <ejb-ref-name>
    oc4j-example/EjbSynchronizations/local
  </ejb-ref-name>
  <ejb-ref-type>Session</ejb-ref-type>
  <local>
    org.jboss.seam.transaction.LocalEjbSynchronizations
  </local>
  <ejb-link>EjbSynchronizations</ejb-link>
</ejb-local-ref>
```

build.xml

Change the default target to archive (we aren't going to cover automatic deployment to OC4J).

Now, lets add in the extra dependencies:

- Hibernate —
 - Copy the jars from hibernate/lib directory in the Seam distribution oc4j-example/lib: cp ../jboss-seam/hibernate/lib/*.jar lib/
 - Alter the build.xml to include them in the ear - add these includes underneath the other libraries being copies:

```
<include name="lib/hibernate-annotations.jar" />
<include name="lib/hibernate-entitymanager.jar" />
<include name="lib/hibernate3.jar" />
<include name="ejb3-peristence.jar" />
<include name="lib/jboss-archive-browsing.jar" />
<include name="lib/jboss-common.jar" />
```

- thirdparty-all.jar — alter the build.xml to include it - add this include:

```
<include name="lib/thirdparty-all.jar" />
```

- antlr-2.7.6.jar — alter the build.xml to include it - add this include:

```
<include name="lib/antlr-*.jar" />
```

- As we are using Drools to provide Seam Security rules, we need to add in Eclipse JDT compiler (you don't need this on JBoss AS; again this is due to OC4J's classloading):

-

```
cp ../jboss-seam/seam-gen/lib/org.eclipse.jdt.core*.jar lib/
```

- Alter the build.xml to include them in the ear:

```
<include name="lib/org.eclipse.jdt.core*.jar" />
```

You should end up with something like:

```
<fileset dir="${basedir}">
  <!-- other libraries added by seam-gen -->
  <include name="lib/hibernate-annotations.jar" />
  <include name="lib/hibernate-entitymanager.jar" />
  <include name="lib/hibernate3.jar" />
  <include name="lib/jboss-archive-browsing.jar" />
  <include name="lib/jboss-common.jar" />
  <include name="lib/thirdparty-all.jar" />
  <include name="lib/antlr-*.jar" />
  <include name="lib/org.eclipse.jdt.core*.jar" />
</fileset>
```

Finally, let's link our User entity into Seam Security (we have a User table with a username column and a password column). We're going to make our authenticator a Stateless Session Bean (OC4J is a EJB3 container after all!):

1.

- Add the `@Stateless` annotation.
- Rename the class to `AuthenticatorAction`
- Create an interface called `Authenticator` which `AuthenticatorAction` implements (EJB3 requires session beans to have a local interface). Annotate the interface with `@Local`, and add a single method with same signature as the `authenticate` in `AuthenticatorAction`.

```
@Name("authenticator") @Stateless public class
    AuthenticatorAction implements Authenticator {
```

```
@Local public interface Authenticator {
    public boolean authenticate();
}
```

2. Use `@PersistenceContext` to inject an `EntityManager`:

```
@PersistenceContext private EntityManager entityManager;
```

3. Implement `authenticate`:

```
public boolean authenticate() {
    List <User> users = entityManager.createQuery("select u from User u where
    u.username = #{identity.username} and
    u.password = #{identity.password}") .getResultList();
    if (users.size() == 1) {
        identity.addRole("admin");
        return true;
    } else {
        return false;
    }
}
```

4. And then add the EJB3 reference to `web.xml`:

```
<ejb-local-ref>
  <ejb-ref-name>
    oc4j-example/AuthenticatorAction/local
  </ejb-ref-name>
  <ejb-ref-type>Session</ejb-ref-type>
  <local>
    org.jboss.seam.tutorial.oc4j.action.Authenticator
  </local>
  <ejb-link>AuthenticatorAction</ejb-link>
</ejb-local-ref>
```

Now you can go on and customize your application.

26.3.1. OC4J Deployment Descriptors for the seam-gen'd application

To deploy your application use the deployment instructions above in conjunction with these deployment descriptors:

`$ORACLE_HOME/j2ee/home/config/data-sources.xml`

```
<managed-data-source
  connection-pool-name="oc4j-example-connection-pool"
  jndi-name="jdbc/__oc4jExample"
  name="oc4j-example-managed-data-source" />
<connection-pool
  name="oc4j-example-connection-pool">
  <connection-factory
    factory-class="com.mysql.jdbc.Driver"
    user="username"
    password="password"
    url="jdbc:mysql:///oc4j" />
  </connection-factory>
</connection-pool>
```

`$ORACLE_HOME/j2ee/home/config/server.xml`

```
<application name="oc4j-example"
  path="../../../home/applications/oc4j-example.ear"
  parent="default"
  start="true" />
```

`$ORACLE_HOME/j2ee/home/config/default-web-site.xml`

```
<web-app application="oc4j-example"
  name="oc4j-example"
  load-on-startup="true"
  root="/oc4j-example" />
```

Chapter 27. Seam annotations

When you write a Seam application, you'll use a lot of annotations. Seam lets you use annotations to achieve a declarative style of programming. Most of the annotations you'll use are defined by the EJB 3.0 specification. The annotations for data validation are defined by the Hibernate Validator package. Finally, Seam defines its own set of annotations, which we'll describe in this chapter.

All of these annotations are defined in the package `org.jboss.seam.annotations`.

27.1. Annotations for component definition

The first group of annotations lets you define a Seam component. These annotations appear on the component class.

`@Name`

```
@Name("componentName")
```

Defines the Seam component name for a class. This annotation is required for all Seam components.

`@Scope`

```
@Scope(ScopeType.CONVERSATION)
```

Defines the default context of the component. The possible values are defined by the `ScopeType` enumeration: `EVENT`, `PAGE`, `CONVERSATION`, `SESSION`, `BUSINESS_PROCESS`, `APPLICATION`, `STATELESS`.

When no scope is explicitly specified, the default depends upon the component type. For stateless session beans, the default is `STATELESS`. For entity beans and stateful session beans, the default is `CONVERSATION`. For JavaBeans, the default is `EVENT`.

`@Role`

```
@Role(name="roleName", scope=ScopeType.SESSION)
```

Allows a Seam component to be bound to multiple contexts variables. The `@Name/@Scope` annotations define a "default role". Each `@Role` annotation defines an additional role.

- `name` — the context variable name.
- `scope` — the context variable scope. When no scope is explicitly specified, the default depends upon the component type, as above.

`@Roles`

```
@Roles({
    @Role(name="user", scope=ScopeType.CONVERSATION),
    @Role(name="currentUser", scope=ScopeType.SESSION)
})
```

Allows specification of multiple additional roles.

@BypassInterceptors

```
@BypassInterceptors
```

Disables Seam all interceptors on a particular component or method of a component.

@JndiName

```
@JndiName("my/jndi/name")
```

Specifies the JNDI name that Seam will use to look up the EJB component. If no JNDI name is explicitly specified, Seam will use the JNDI pattern specified by `org.jboss.seam.core.init.jndiPattern`.

@Conversational

```
@Conversational
```

Specifies that a conversation scope component is conversational, meaning that no method of the component may be called unless a long-running conversation is active.

@Startup

```
@Scope(APPLICATION) @Startup(depends="org.jboss.seam.bpm.jbpm")
```

Specifies that an application scope component is started immediately at initialization time. This is mainly used for certain built-in components that bootstrap critical infrastructure such as JNDI, datasources, etc.

```
@Scope(SESSION) @Startup
```

Specifies that a session scope component is started immediately at session creation time.

- `depends` — specifies that the named components must be started first, if they are installed.

@Install

```
@Install(false)
```

Specifies whether or not a component should be installed by default. The lack of an `@Install` annotation indicates a component should be installed.

```
@Install(dependencies="org.jboss.seam.bpm.jbpm")
```

Specifies that a component should only be stalled if the components listed as dependencies are also installed.

```
@Install(genericDependencies=ManagedQueueSender.class)
```

Specifies that a component should only be installed if a component that is implemented by a certain class is installed. This is useful when the dependency doesn't have a single well-known name.

```
@Install(classDependencies="org.hibernate.Session")
```

Specifies that a component should only be installed if the named class is in the classpath.

```
@Install(precedence=BUILT_IN)
```

Specifies the precedence of the component. If multiple components with the same name exist, the one with the higher precedence will be installed. The defined precedence values are (in ascending order):

- `BUILT_IN` — Precedence of all built-in Seam components
- `FRAMEWORK` — Precedence to use for components of frameworks which extend Seam
- `APPLICATION` — Precedence of application components (the default precedence)
- `DEPLOYMENT` — Precedence to use for components which override application components in a particular deployment
- `MOCK` — Precedence for mock objects used in testing

`@Synchronized`

```
@Synchronized(timeout=1000)
```

Specifies that a component is accessed concurrently by multiple clients, and that Seam should serialize requests. If a request is not able to obtain its lock on the component in the given timeout period, an exception will be raised.

`@ReadOnly`

```
@ReadOnly
```

Specifies that a JavaBean component or component method does not require state replication at the end of the invocation.

`@AutoCreate`

```
@AutoCreate
```

Specifies that a component will be automatically created, even if the client does not specify `create=true`.

27.2. Annotations for bijection

The next two annotations control bijection. These attributes occur on component instance variables or property accessor methods.

`@In`

```
@In
```

Specifies that a component attribute is to be injected from a context variable at the beginning of each component invocation. If the context variable is null, an exception will be thrown.

```
@In(required=false)
```

Specifies that a component attribute is to be injected from a context variable at the beginning of each com-

ponent invocation. The context variable may be null.

```
@In(create=true)
```

Specifies that a component attribute is to be injected from a context variable at the beginning of each component invocation. If the context variable is null, an instance of the component is instantiated by Seam.

```
@In(value="contextVariableName")
```

Specifies the name of the context variable explicitly, instead of using the annotated instance variable name.

```
@In(value="#{customer.addresses['shipping']}")
```

Specifies that a component attribute is to be injected by evaluating a JSF EL expression at the beginning of each component invocation.

- `value` — specifies the name of the context variable. Default to the name of the component attribute. Alternatively, specifies a JSF EL expression, surrounded by `#{...}`.
- `create` — specifies that Seam should instantiate the component with the same name as the context variable if the context variable is undefined (null) in all contexts. Default to false.
- `required` — specifies Seam should throw an exception if the context variable is undefined in all contexts.

@Out

```
@Out
```

Specifies that a component attribute that is a Seam component is to be outjected to its context variable at the end of the invocation. If the attribute is null, an exception is thrown.

```
@Out(required=false)
```

Specifies that a component attribute that is a Seam component is to be outjected to its context variable at the end of the invocation. The attribute may be null.

```
@Out(scope=ScopeType.SESSION)
```

Specifies that a component attribute that is *not* a Seam component type is to be outjected to a specific scope at the end of the invocation.

Alternatively, if no scope is explicitly specified, the scope of the component with the @Out attribute is used (or the EVENT scope if the component is stateless).

```
@Out(value="contextVariableName")
```

Specifies the name of the context variable explicitly, instead of using the annotated instance variable name.

- `value` — specifies the name of the context variable. Default to the name of the component attribute.
- `required` — specifies Seam should throw an exception if the component attribute is null during outjection.

Note that it is quite common for these annotations to occur together, for example:

```
@In(create=true) @Out private User currentUser;
```

The next annotation supports the *manager component* pattern, where a Seam component manages the life-cycle of an instance of some other class that is to be injected. It appears on a component getter method.

@Unwrap

```
@Unwrap
```

Specifies that the object returned by the annotated getter method is the thing that is injected instead of the component instance itself.

The next annotation supports the *factory component* pattern, where a Seam component is responsible for initializing the value of a context variable. This is especially useful for initializing any state needed for rendering the response to a non-faces request. It appears on a component method.

@Factory

```
@Factory("processInstance") public void createProcessInstance() { ... }
```

Specifies that the method of the component is used to initialize the value of the named context variable, when the context variable has no value. This style is used with methods that return void.

```
@Factory("processInstance", scope=CONVERSATION) public ProcessInstance createProcessInstance() { ... }
```

Specifies that the method returns a value that Seam should use to initialize the value of the named context variable, when the context variable has no value. This style is used with methods that return a value. If no scope is explicitly specified, the scope of the component with the @Factory method is used (unless the component is stateless, in which case the EVENT context is used).

- `value` — specifies the name of the context variable. If the method is a getter method, default to the JavaBeans property name.
- `scope` — specifies the scope that Seam should bind the returned value to. Only meaningful for factory methods which return a value.
- `autoCreate` — specifies that this factory method should be automatically called whenever the variable is asked for, even if @In does not specify `create=true`.

This annotation lets you inject a Log:

@Logger

```
@Logger("categoryName")
```

Specifies that a component field is to be injected with an instance of `org.jboss.seam.log.Log`. For entity beans, the field must be declared as static.

- `value` — specifies the name of the log category. Default to the name of the component class.

The last annotation lets you inject a request parameter value:

`@RequestParam`

```
@RequestParam("parameterName")
```

Specifies that a component attribute is to be injected with the value of a request parameter. Basic type conversions are performed automatically.

- `value` — specifies the name of the request parameter. Default to the name of the component attribute.

27.3. Annotations for component lifecycle methods

These annotations allow a component to react to its own lifecycle events. They occur on methods of the component. There may be only one of each per component class.

`@Create`

```
@Create
```

Specifies that the method should be called when an instance of the component is instantiated by Seam. Note that create methods are only supported for JavaBeans and stateful session beans.

`@Destroy`

```
@Destroy
```

Specifies that the method should be called when the context ends and its context variables are destroyed. Note that destroy methods are only supported for JavaBeans and stateful session beans.

Destroy methods should be used only for cleanup. *Seam catches, logs and swallows any exception that propagates out of a destroy method.*

`@Observer`

```
@Observer("somethingChanged")
```

Specifies that the method should be called when a component-driven event of the specified type occurs.

```
@Observer(value="somethingChanged", create=false)
```

Specifies that the method should be called when an event of the specified type occurs but that an instance should not be created if one doesn't exist. If an instance does not exist and `create` is false, the event will not be observed. The default value for `create` is true.

27.4. Annotations for context demarcation

These annotations provide declarative conversation demarcation. They appear on methods of Seam components, usually action listener methods.

Every web request has a conversation context associated with it. Most of these conversations end at the end of the request. If you want a conversation that span multiple requests, you must "promote" the current conversation to a *long-running conversation* by calling a method marked with `@Begin`.

`@Begin`

```
@Begin
```

Specifies that a long-running conversation begins when this method returns a non-null outcome without exception.

```
@Begin(join=true)
```

Specifies that if a long-running conversation is already in progress, the conversation context is simply propagated.

```
@Begin(nested=true)
```

Specifies that if a long-running conversation is already in progress, a new *nested* conversation context begins. The nested conversation will end when the next `@End` is encountered, and the outer conversation will resume. It is perfectly legal for multiple nested conversations to exist concurrently in the same outer conversation.

```
@Begin(pageflow="process definition name")
```

Specifies a jBPM process definition name that defines the pageflow for this conversation.

```
@Begin(flushMode=FlushModeType.MANUAL)
```

Specify the flush mode of any Seam-managed persistence contexts. `flushMode=FlushModeType.MANUAL` supports the use of *atomic conversations* where all write operations are queued in the conversation context until an explicit call to `flush()` (which usually occurs at the end of the conversation).

- `join` — determines the behavior when a long-running conversation is already in progress. If `true`, the context is propagated. If `false`, an exception is thrown. Default to `false`. This setting is ignored when `nested=true` is specified
- `nested` — specifies that a nested conversation should be started if a long-running conversation is already in progress.
- `flushMode` — set the flush mode of any Seam-managed Hibernate sessions or JPA persistence contexts that are created during this conversation.
- `pageflow` — a process definition name of a jBPM process definition deployed via `org.jboss.seam.bpm.jbpm.pageflowDefinitions`.

`@End`

```
@End
```

Specifies that a long-running conversation ends when this method returns a non-null outcome without exception.

- `beforeRedirect` — by default, the conversation will not actually be destroyed until after any redirect has occurred. Setting `beforeRedirect=true` specifies that the conversation should be destroyed at the end of the current request, and that the redirect will be processed in a new temporary conversation context.

@StartTask

@StartTask

"Starts" a jBPM task. Specifies that a long-running conversation begins when this method returns a non-null outcome without exception. This conversation is associated with the jBPM task specified in the named request parameter. Within the context of this conversation, a business process context is also defined, for the business process instance of the task instance.

The jBPM `TaskInstance` will be available in a request context variable named `taskInstance`. The jBPM `ProcessInstance` will be available in a request context variable named `processInstance`. (Of course, these objects are available for injection via `@In`.)

- `taskIdParameter` — the name of a request parameter which holds the id of the task. Default to "taskId", which is also the default used by the Seam `taskList` JSF component.
- `flushMode` — set the flush mode of any Seam-managed Hibernate sessions or JPA persistence contexts that are created during this conversation.

@BeginTask

@BeginTask

Resumes work on an incomplete jBPM task. Specifies that a long-running conversation begins when this method returns a non-null outcome without exception. This conversation is associated with the jBPM task specified in the named request parameter. Within the context of this conversation, a business process context is also defined, for the business process instance of the task instance.

The jBPM `org.jbpm.taskmgmt.exe.TaskInstance` will be available in a request context variable named `taskInstance`. The jBPM `org.jbpm.graph.exe.ProcessInstance` will be available in a request context variable named `processInstance`.

- `taskIdParameter` — the name of a request parameter which holds the id of the task. Default to "taskId", which is also the default used by the Seam `taskList` JSF component.
- `flushMode` — set the flush mode of any Seam-managed Hibernate sessions or JPA persistence contexts that are created during this conversation.

@EndTask

@EndTask

"Ends" a jBPM task. Specifies that a long-running conversation ends when this method returns a non-null outcome, and that the current task is complete. Triggers a jBPM transition. The actual transition triggered will be the default transition unless the application has called `Transition.setName()` on the built-in component named `transition`.

```
@EndTask(transition="transitionName")
```

Triggers the given jBPM transition.

- `transition` — the name of the jBPM transition to be triggered when ending the task. Defaults to the default transition.
- `beforeRedirect` — by default, the conversation will not actually be destroyed until after any redirect has occurred. Setting `beforeRedirect=true` specifies that the conversation should be destroyed at the end of the current request, and that the redirect will be processed in a new temporary conversation context.

@CreateProcess

```
@CreateProcess(definition="process definition name")
```

Creates a new jBPM process instance when the method returns a non-null outcome without exception. The `ProcessInstance` object will be available in a context variable named `processInstance`.

- `definition` — the name of the jBPM process definition deployed via `org.jboss.seam.bpm.jbpm.processDefinitions`.

@ResumeProcess

```
@ResumeProcess(processIdParameter="processId")
```

Re-enters the scope of an existing jBPM process instance when the method returns a non-null outcome without exception. The `ProcessInstance` object will be available in a context variable named `processInstance`.

- `processIdParameter` — the name a request parameter holding the process id. Default to `"processId"`.

@Transition

```
@Transition("cancel")
```

Marks a method as signalling a transition in the current jBPM process instance whenever the method returns a non-null result.

27.5. Annotations for use with Seam JavaBean components in a J2EE environment

Seam provides an annotation that lets you force a rollback of the JTA transaction for certain action listener outcomes.

@Transactional

```
@Transactional
```

Specifies that a JavaBean component should have a similar transactional behavior to the default behavior of a session bean component. ie. method invocations should take place in a transaction, and if no transaction exists when the method is called, a transaction will be started just for that method. This annotation may be

applied at either class or method level. *Do not use this annotations on EJB 3.0 components, use `@TransactionAttribute`!*

`@ApplicationException`

```
@Transactional
```

TDB

`@Interceptors`

```
@Transactional
```

TDB

These annotations are mostly useful for JavaBean Seam components. If you use EJB 3.0 components, you should use the standard `@TransactionAttribute` annotation.

27.6. Annotations for exceptions

These annotations let you specify how Seam should handle an exception that propagates out of a Seam component.

`@Redirect`

```
@Redirect(viewId="error.jsp")
```

Specifies that the annotated exception causes a browser redirect to a specified view id.

- `viewId` — specifies the JSF view id to redirect to. You can use EL here.
- `message` — a message to be displayed, default to the exception message.
- `end` — specifies that the long-running conversation should end, default to `false`.

`@HttpError`

```
@HttpError(errorCode=404)
```

Specifies that the annotated exception causes a HTTP error to be sent.

- `errorCode` — the HTTP error code, default to 500.
- `message` — a message to be sent with the HTTP error, default to the exception message.
- `end` — specifies that the long-running conversation should end, default to `false`.

27.7. Annotations for Seam Remoting

Seam Remoting requires that the local interface of a session bean be annotated with the following annotation:

@WebRemote

```
@WebRemote(exclude="path.to.exclude")
```

Indicates that the annotated method may be called from client-side JavaScript. The `exclude` property is optional and allows objects to be excluded from the result's object graph (see the Remoting chapter for more details).

27.8. Annotations for Seam interceptors

The following annotations appear on Seam interceptor classes.

Please refer to the documentation for the EJB 3.0 specification for information about the annotations required for EJB interceptor definition.

@Interceptor

```
@Interceptor(stateless=true)
```

Specifies that this interceptor is stateless and Seam may optimize replication.

```
@Interceptor(type=CLIENT)
```

Specifies that this interceptor is a "client-side" interceptor that is called before the EJB container.

```
@Interceptor(around={SomeInterceptor.class, OtherInterceptor.class})
```

Specifies that this interceptor is positioned higher in the stack than the given interceptors.

```
@Interceptor(within={SomeInterceptor.class, OtherInterceptor.class})
```

Specifies that this interceptor is positioned deeper in the stack than the given interceptors.

27.9. Annotations for asynchronicity

The following annotations are used to declare an asynchronous method, for example:

```
@Asynchronous public void scheduleAlert(Alert alert, @Expiration Date date) { ... }
```

```
@Asynchronous public Timer scheduleAlerts(Alert alert,
                                           @Expiration Date date,
                                           @IntervalDuration long interval) { ... }
```

@Asynchronous

```
@Asynchronous
```

Specifies that the method call is processed asynchronously.

@Duration

```
@Duration
```

Specifies that a parameter of the asynchronous call is the duration before the call is processed (or first processed for recurring calls).

@Expiration

```
@Expiration
```

Specifies that a parameter of the asynchronous call is the datetime at which the call is processed (or first processed for recurring calls).

@IntervalDuration

```
@IntervalDuration
```

Specifies that an asynchronous method call recurs, and that the annotated parameter is duration between recurrences.

27.10. Annotations for use with JSF

The following annotations make working with JSF easier.

@Converter

Allows a Seam component to act as a JSF converter. The annotated class must be a Seam component, and must implement `javax.faces.convert.Converter`.

- `id` — the JSF converter id. Defaults to the component name.
- `forClass` — if specified, register this component as the default converter for a type.

@Validator

Allows a Seam component to act as a JSF validator. The annotated class must be a Seam component, and must implement `javax.faces.validator.Validator`.

- `id` — the JSF validator id. Defaults to the component name.

27.10.1. Annotations for use with dataTable

The following annotations make it easy to implement clickable lists backed by a stateful session bean. They appear on attributes.

@DataModel

```
@DataModel("variableName")
```

Outjects a property of type `List`, `Map`, `Set` or `Object[]` as a JSF `DataModel` into the scope of the owning component (or the `EVENT` scope if the owning component is `STATELESS`). In the case of `Map`, each row of the `DataModel` is a `Map.Entry`.

- `value` — name of the conversation context variable. Default to the attribute name.
- `scope` — if `scope=ScopeType.PAGE` is explicitly specified, the `DataModel` will be kept in the `PAGE` context.

`@DataModelSelection`

```
@DataModelSelection
```

Injects the selected value from the JSF `DataModel` (this is the element of the underlying collection, or the map value). If only one `@DataModel` attribute is defined for a component, the selected value from that `DataModel` will be injected. Otherwise, the component name of each `@DataModel` must be specified in the `value` attribute for each `@DataModelSelection`.

If `PAGE` scope is specified on the associated `@DataModel`, then, in addition to the `DataModel Selection` being injected, the associated `DataModel` will also be injected. In this case, if the property annotated with `@DataModel` is a getter method, then a setter method for the property must also be part of the Business API of the containing Seam Component.

- `value` — name of the conversation context variable. Not needed if there is exactly one `@DataModel` in the component.

`@DataModelSelectionIndex`

```
@DataModelSelectionIndex
```

Exposes the selection index of the JSF `DataModel` as an attribute of the component (this is the row number of the underlying collection, or the map key). If only one `@DataModel` attribute is defined for a component, the selected value from that `DataModel` will be injected. Otherwise, the component name of each `@DataModel` must be specified in the `value` attribute for each `@DataModelSelectionIndex`.

- `value` — name of the conversation context variable. Not needed if there is exactly one `@DataModel` in the component.

27.11. Meta-annotations for databinding

These meta-annotations make it possible to implement similar functionality to `@DataModel` and `@DataModelSelection` for other datastructures apart from lists.

`@DataBinderClass`

```
@DataBinderClass(DataModelBinder.class)
```

Specifies that an annotation is a databinding annotation.

`@DataSelectorClass`

```
@DataSelectorClass(DataModelSelector.class)
```

Specifies that an annotation is a dataselection annotation.

27.12. Annotations for packaging

This annotation provides a mechanism for declaring information about a set of components that are packaged together. It can be applied to any Java package.

@Namespace

```
@Namespace(value="http://jboss.com/products/seam/example/seampay")
```

Specifies that components in the current package are associated with the given namespace. The declared namespace can be used as an XML namespace in a `components.xml` file to simplify application configuration.

```
@Namespace(value="http://jboss.com/products/seam/core", prefix="org.jboss.seam.core")
```

Specifies a namespace to associate with a given package. Additionally, it specifies a component name prefix to be applied to component names specified in the XML file. For example, an XML element named `init` that is associated with this namespace would be understood to actually refer to a component named `org.jboss.seam.core.init`.

27.13. Annotations for integrating with the servlet container

These annotations allow you to integrate your Seam components with the servlet container.

@Filter

Use the Seam component (which implements `javax.servlet.Filter`) annotated with `@Filter` as a servlet filter. It will be executed by Seam's master filter.

•

```
@Filter(around={"seamComponent", "otherSeamComponent"})
```

Specifies that this filter is positioned higher in the stack than the given filters.

•

```
@Filter(within={"seamComponent", "otherSeamComponent"})
```

Specifies that this filter is positioned deeper in the stack than the given filters.

Chapter 28. Built-in Seam components

This chapter describes Seam's built-in components, and their configuration properties. The built-in components will be created even if they are not listed in your `components.xml` file, but if you need to override default properties or specify more than one component of a certain type, `components.xml` is used.

Note that you can replace any of the built in components with your own implementations simply by specifying the name of one of the built in components on your own class using `@Name`.

Note also that even though all the built in components use a qualified name, most of them are aliased to unqualified names by default. These aliases specify `auto-create="true"`, so you do not need to use `create=true` when injecting built-in components by their unqualified name.

28.1. Context injection components

The first set of built in components exist purely to support injection of various contextual objects. For example, the following component instance variable would have the Seam session context object injected:

```
@In private Context sessionContext;
```

`org.jboss.seam.core.contexts`

Component that provides access to Seam Context objects, for example
`org.jboss.seam.core.contexts.sessionContext['user']`.

`org.jboss.seam.faces.facesContext`

Manager component for the `FacesContext` context object (not a true Seam context)

All of these components are always installed.

28.2. Utility components

These components are merely useful.

`org.jboss.seam.faces.facesMessages`

Allows faces success messages to propagate across a browser redirect.

- `add(FacesMessage facesMessage)` — add a faces message, which will be displayed during the next render response phase that occurs in the current conversation.
- `add(String messageTemplate)` — add a faces message, rendered from the given message template which may contain EL expressions.
- `add(Severity severity, String messageTemplate)` — add a faces message, rendered from the given message template which may contain EL expressions.
- `addFromResourceBundle(String key)` — add a faces message, rendered from a message template defined in the Seam resource bundle which may contain EL expressions.
- `addFromResourceBundle(Severity severity, String key)` — add a faces message, rendered from a

message template defined in the Seam resource bundle which may contain EL expressions.

- `clear()` — clear all messages.

`org.jboss.seam.faces.redirect`

A convenient API for performing redirects with parameters (this is especially useful for bookmarkable search results screens).

- `redirect.viewId` — the JSF view id to redirect to.
- `redirect.conversationPropagationEnabled` — determines whether the conversation will propagate across the redirect.
- `redirect.parameters` — a map of request parameter name to value, to be passed in the redirect request.
- `execute()` — perform the redirect immediately.
- `captureCurrentRequest()` — stores the view id and request parameters of the current GET request (in the conversation context), for later use by calling `execute()`.

`org.jboss.seam.faces.httpError`

A convenient API for sending HTTP errors.

`org.jboss.seam.core.events`

An API for raising events that can be observed via `@Observer` methods, or method bindings in `components.xml`.

- `raiseEvent(String type)` — raise an event of a particular type and distribute to all observers.
- `raiseAsynchronousEvent(String type)` — raise an event to be processed asynchronously by the EJB3 timer service.
- `raiseTimedEvent(String type, ....)` — schedule an event to be processed asynchronously by the EJB3 timer service.
- `addListener(String type, String methodBinding)` — add an observer for a particular event type.

`org.jboss.seam.core.interpolator`

An API for interpolating the values of JSF EL expressions in Strings.

- `interpolate(String template)` — scan the template for JSF EL expressions of the form `#{...}` and replace them with their evaluated values.

`org.jboss.seam.core.expressions`

An API for creating value and method bindings.

- `createValueBinding(String expression)` — create a value binding object.
- `createMethodBinding(String expression)` — create a method binding object.

`org.jboss.seam.core.pojoCache`

Manager component for a JBoss Cache `PojoCache` instance.

- `pojoCache.cfgResourceName` — the name of the configuration file. Default to `treecache.xml`.

All of these components are always installed.

28.3. Components for internationalization and themes

The next group of components make it easy to build internationalized user interfaces using Seam.

`org.jboss.seam.core.locale`

The Seam locale.

`org.jboss.seam.international.timezone`

The Seam timezone. The timezone is session scoped.

`org.jboss.seam.core.resourceBundle`

The Seam resource bundle. The resource bundle is stateless. The Seam resource bundle performs a depth-first search for keys in a list of Java resource bundles.

`org.jboss.seam.core.resourceLoader`

The resource loader provides access to application resources and resource bundles.

- `resourceLoader.bundleNames` — the names of the Java resource bundles to search when the Seam resource bundle is used. Default to `messages`.

`org.jboss.seam.international.localeSelector`

Supports selection of the locale either at configuration time, or by the user at runtime.

- `select()` — select the specified locale.
- `localeSelector.locale` — the actual `java.util.Locale`.
- `localeSelector.localeString` — the stringified representation of the locale.
- `localeSelector.language` — the language for the specified locale.
- `localeSelector.country` — the country for the specified locale.
- `localeSelector.variant` — the variant for the specified locale.
- `localeSelector.supportedLocales` — a list of `SelectItems` representing the supported locales listed in `jsf-config.xml`.
- `localeSelector.cookieEnabled` — specifies that the locale selection should be persisted via a cookie.

`org.jboss.seam.international.timezoneSelector`

Supports selection of the timezone either at configuration time, or by the user at runtime.

- `select()` — select the specified locale.
- `timezoneSelector.timezone` — the actual `java.util.TimeZone`.
- `timezoneSelector.timeZoneId` — the stringified representation of the timezone.

- `timezoneSelector.cookieEnabled` — specifies that the timezone selection should be persisted via a cookie.

`org.jboss.seam.international.messages`

A map containing internationalized messages rendered from message templates defined in the Seam resource bundle.

`org.jboss.seam.theme.themeSelector`

Supports selection of the theme either at configuration time, or by the user at runtime.

- `select()` — select the specified theme.
- `theme.availableThemes` — the list of defined themes.
- `themeSelector.theme` — the selected theme.
- `themeSelector.themes` — a list of `SelectItems` representing the defined themes.
- `themeSelector.cookieEnabled` — specifies that the theme selection should be persisted via a cookie.

`org.jboss.seam.theme.theme`

A map containing theme entries.

All of these components are always installed.

28.4. Components for controlling conversations

The next group of components allow control of conversations by the application or user interface.

`org.jboss.seam.core.conversation`

API for application control of attributes of the current Seam conversation.

- `getId()` — returns the current conversation id
- `isNested()` — is the current conversation a nested conversation?
- `isLongRunning()` — is the current conversation a long-running conversation?
- `getId()` — returns the current conversation id
- `getParentId()` — returns the conversation id of the parent conversation
- `getRootId()` — returns the conversation id of the root conversation
- `setTimeout(int timeout)` — sets the timeout for the current conversation
- `setViewId(String outcome)` — sets the view id to be used when switching back to the current conversation from the conversation switcher, conversation list, or breadcrumbs.
- `setDescription(String description)` — sets the description of the current conversation to be displayed in the conversation switcher, conversation list, or breadcrumbs.
- `redirect()` — redirect to the last well-defined view id for this conversation (useful after login chal-

lenges).

- `leave()` — exit the scope of this conversation, without actually ending the conversation.
- `begin()` — begin a long-running conversation (equivalent to `@Begin`).
- `beginPageflow(String pageflowName)` — begin a long-running conversation with a pageflow (equivalent to `@Begin(pageflow="...")`).
- `end()` — end a long-running conversation (equivalent to `@End`).
- `pop()` — pop the conversation stack, returning to the parent conversation.
- `root()` — return to the root conversation of the conversation stack.
- `changeFlushMode(FlushModeType flushMode)` — change the flush mode of the conversation.

`org.jboss.seam.core.conversationList`
Manager component for the conversation list.

`org.jboss.seam.core.conversationStack`
Manager component for the conversation stack (breadcrumbs).

`org.jboss.seam.faces.switcher`
The conversation switcher.

All of these components are always installed.

28.5. jBPM-related components

These components are for use with jBPM.

`org.jboss.seam.pageflow.pageflow`
API control of Seam pageflows.

- `isInProcess()` — returns true if there is currently a pageflow in process
- `getProcessInstance()` — returns jBPM `ProcessInstance` for the current pageflow
- `begin(String pageflowName)` — begin a pageflow in the context of the current conversation
- `reposition(String nodeName)` — reposition the current pageflow to a particular node

`org.jboss.seam.bpm.actor`
API for application control of attributes of the jBPM actor associated with the current session.

- `setId(String actorId)` — sets the jBPM actor id of the current user.
- `getGroupActorIds()` — returns a Set to which jBPM actor ids for the current users groups may be added.

`org.jboss.seam.bpm.transition`
API for application control of the jBPM transition for the current task.

- `setName(String transitionName)` — sets the jBPM transition name to be used when the current task is ended via `@EndTask`.

`org.jboss.seam.bpm.businessProcess`

API for programmatic control of the association between the conversation and business process.

- `businessProcess.taskId` — the id of the task associated with the current conversation.
- `businessProcess.processId` — the id of the process associated with the current conversation.
- `businessProcess.hasCurrentTask()` — is a task instance associated with the current conversation?
- `businessProcess.hasCurrentProcess()` — is a process instance associated with the current conversation.
- `createProcess(String name)` — create an instance of the named process definition and associate it with the current conversation.
- `startTask()` — start the task associated with the current conversation.
- `endTask(String transitionName)` — end the task associated with the current conversation.
- `resumeTask(Long id)` — associate the task with the given id with the current conversation.
- `resumeProcess(Long id)` — associate the process with the given id with the current conversation.
- `transition(String transitionName)` — trigger the transition.

`org.jboss.seam.bpm.taskInstance`

Manager component for the jBPM TaskInstance.

`org.jboss.seam.bpm.processInstance`

Manager component for the jBPM ProcessInstance.

`org.jboss.seam.bpm.jbpmContext`

Manager component for an event-scoped JbpmContext.

`org.jboss.seam.bpm.taskInstanceList`

Manager component for the jBPM task list.

`org.jboss.seam.bpm.pooledTaskInstanceList`

Manager component for the jBPM pooled task list.

`org.jboss.seam.bpm.taskInstanceListForType`

Manager component for the jBPM task lists.

`org.jboss.seam.bpm.pooledTask`

Action handler for pooled task assignment.

`org.jboss.seam.bpm.processInstanceFinder`

Manager for the process instance task list.

`org.jboss.seam.bpm.processInstanceList`

The process instance task list.

All of these components are installed whenever the component `org.jboss.seam.bpm.jbpm` is installed.

28.6. Security-related components

These components relate to web-tier security.

`org.jboss.seam.web.userPrincipal`

Manager component for the current user Principal.

`org.jboss.seam.web.isUserInRole`

Allows JSF pages to choose to render a control, depending upon the roles available to the current principal.

`<h:commandButton value="edit" rendered="#{isUserInRole['admin']}" />`.

28.7. JMS-related components

These components are for use with managed TopicPublishers and QueueSenders (see below).

`org.jboss.seam.jms.queueSession`

Manager component for a JMS QueueSession .

`org.jboss.seam.jms.topicSession`

Manager component for a JMS TopicSession .

28.8. Mail-related components

These components are for use with Seam's Email support

`org.jboss.seam.mail.mailSession`

Manager component for a JavaMail Session. The session can be either looked up in the JNDI context (by setting the `sessionJndiName` property) or it can be created from the configuration options in which case the host is mandatory.

- `org.jboss.seam.mail.mailSession.host` — the hostname of the SMTP server to use
- `org.jboss.seam.mail.mailSession.port` — the port of the SMTP server to use
- `org.jboss.seam.mail.mailSession.username` — the username to use to connect to the SMTP server.
- `org.jboss.seam.mail.mailSession.password` — the password to use to connect to the SMTP server
- `org.jboss.seam.mail.mailSession.debug` — enable JavaMail debugging (very verbose)
- `org.jboss.seam.mail.mailSession.ssl` — enable SSL connection to SMTP (will default to port 465)
- `org.jboss.seam.mail.mailSession.tls` — by default true, enable TLS support in the mail session
- `org.jboss.seam.mail.mailSession.sessionJndiName` — name under which a `javax.mail.Session` is bound to JNDI. If supplied, all other properties will be ignored.

28.9. Infrastructural components

These components provide critical platform infrastructure. You can install a component which isn't installed by default by setting `install="true"` on the component in `components.xml`.

`org.jboss.seam.core.init`

Initialization settings for Seam. Always installed.

- `org.jboss.seam.core.init.jndiPattern` — the JNDI pattern used for looking up session beans
- `org.jboss.seam.core.init.debug` — enable Seam debug mode
- `org.jboss.seam.core.init.clientSideConversations` — if set to `true`, Seam will save conversation context variables in the client instead of in the `HttpSession`.
- `org.jboss.seam.core.init.userTransactionName` — the JNDI name to use when looking up the JTA `UserTransaction` object.

`org.jboss.seam.core.manager`

Internal component for Seam page and conversation context management. Always installed.

- `org.jboss.seam.core.manager.conversationTimeout` — the conversation context timeout in milliseconds.
- `org.jboss.seam.core.manager.concurrentRequestTimeout` — maximum wait time for a thread attempting to gain a lock on the long-running conversation context.
- `org.jboss.seam.core.manager.conversationIdParameter` — the request parameter used to propagate the conversation id, default to `conversationId`.
- `org.jboss.seam.core.manager.conversationIsLongRunningParameter` — the request parameter used to propagate information about whether the conversation is long-running, default to `conversationIsLongRunning`.

`org.jboss.seam.navigation.pages`

Internal component for Seam workspace management. Always installed.

- `org.jboss.seam.navigation.pages.noConversationViewId` — global setting for the view id to redirect to when a conversation entry is not found on the server side.
- `org.jboss.seam.navigation.pages.loginViewId` — global setting for the view id to redirect to when an unauthenticated user tries to access a protected view.
- `org.jboss.seam.navigation.pages.httpPort` — global setting for the port to use when the http scheme is requested.
- `org.jboss.seam.navigation.pages.httpsPort` — global setting for the port to use when the https scheme is requested.
- `org.jboss.seam.navigation.pages.resources` — a list of resources to search for `pages.xml` style resources. Defaults to `WEB-INF/pages.xml`.

`org.jboss.seam.bpm.jbpm`

Bootstraps a `JbpmConfiguration`. Install as class `org.jboss.seam.bpm.Jbpm`.

- `org.jboss.seam.bpm.jbpm.processDefinitions` — a list of resource names of jPDL files to be used for orchestration of business processes.
- `org.jboss.seam.bpm.jbpm.pageflowDefinitions` — a list of resource names of jPDL files to be used for orchestration of conversation page flows.

`org.jboss.seam.core.conversationEntries`

Internal session-scoped component recording the active long-running conversations between requests.

`org.jboss.seam.faces.facesPage`

Internal page-scoped component recording the conversation context associated with a page.

`org.jboss.seam.persistence.persistenceContexts`

Internal component recording the persistence contexts which were used in the current conversation.

`org.jboss.seam.jms.queueConnection`

Manages a JMS `QueueConnection`. Installed whenever managed `QueueSender` is installed.

- `org.jboss.seam.jms.queueConnection.queueConnectionFactoryJndiName` — the JNDI name of a JMS `QueueConnectionFactory`. Default to `UIL2ConnectionFactory`

`org.jboss.seam.jms.topicConnection`

Manages a JMS `TopicConnection`. Installed whenever managed `TopicPublisher` is installed.

- `org.jboss.seam.jms.topicConnection.topicConnectionFactoryJndiName` — the JNDI name of a JMS `TopicConnectionFactory`. Default to `UIL2ConnectionFactory`

`org.jboss.seam.persistence.persistenceProvider`

Abstraction layer for non-standardized features of JPA provider.

`org.jboss.seam.core.validators`

Caches instances of Hibernate `Validator` `ClassValidator`.

`org.jboss.seam.faces.validation`

Allows the application to determine whether validation failed or was successful.

`org.jboss.seam.debug.introspector`

Support for the Seam Debug Page.

`org.jboss.seam.debug.contexts`

Support for the Seam Debug Page.

`org.jboss.seam.exception.exceptions`

Internal component for exception handling.

`org.jboss.seam.transaction.transaction`

API for controlling transactions and abstracting the underlying transaction management implementation behind a JTA-compatible interface.

`org.jboss.seam.faces.safeActions`

Decides if an action expression in an incoming URL is safe. This is done by checking that the action expression exists in the view.

28.10. Miscellaneous components

These components don't fit into

`org.jboss.seam.async.dispatcher`

Dispatcher stateless session bean for asynchronous methods.

`org.jboss.seam.core.image`

Image manipulation and interrogation.

`org.jboss.seam.core.pojoCache`

Manager component for a PojoCache instance.

`org.jboss.seam.core.uiComponent`

Manages a map of UIComponents keyed by component id.

28.11. Special components

Certain special Seam component classes are installable multiple times under names specified in the Seam configuration. For example, the following lines in `components.xml` install and configure two Seam components:

```
<component name="bookingDatabase"
    class="org.jboss.seam.persistence.ManagedPersistenceContext">
  <property name="persistenceUnitJndiName">java:/comp/emf/bookingPersistence</property>
</component>

<component name="userDatabase"
    class="org.jboss.seam.persistence.ManagedPersistenceContext">
  <property name="persistenceUnitJndiName">java:/comp/emf/userPersistence</property>
</component>
```

The Seam component names are `bookingDatabase` and `userDatabase`.

`<entityManager>`, `org.jboss.seam.persistence.ManagedPersistenceContext`

Manager component for a conversation scoped managed `EntityManager` with an extended persistence context.

- `<entityManager>.entityManagerFactory` — a value binding expression that evaluates to an instance of `EntityManagerFactory`.

`<entityManager>.persistenceUnitJndiName` — the JNDI name of the entity manager factory, default to `java:/<managedPersistenceContext>`.

`<entityManagerFactory>`, `org.jboss.seam.persistence.EntityManagerFactory`

Manages a JPA `EntityManagerFactory`. This is most useful when using JPA outside of an EJB 3.0 supporting environment.

- `entityManagerFactory.persistenceUnitName` — the name of the persistence unit.

See the API JavaDoc for further configuration properties.

`<session>`, `org.jboss.seam.persistence.ManagedSession`

Manager component for a conversation scoped managed Hibernate Session.

- `<session>.sessionFactory` — a value binding expression that evaluates to an instance of `SessionFactory`.

`<session>.sessionFactoryJndiName` — the JNDI name of the session factory, default to `java:/*managedSession>`.

`<sessionFactory>`, `org.jboss.seam.persistence.HibernateSessionFactory`

Manages a Hibernate `SessionFactory`.

- `<sessionFactory>.cfgResourceName` — the path to the configuration file. Default to `hibernate.cfg.xml`.

See the API JavaDoc for further configuration properties.

`<managedQueueSender>`, `org.jboss.seam.jms.ManagedQueueSender`

Manager component for an event scoped managed JMS `QueueSender`.

- `<managedQueueSender>.queueJndiName` — the JNDI name of the JMS queue.

`<managedTopicPublisher>`, `org.jboss.seam.jms.ManagedTopicPublisher`

Manager component for an event scoped managed JMS `TopicPublisher`.

- `<managedTopicPublisher>.topicJndiName` — the JNDI name of the JMS topic.

`<managedWorkingMemory>`, `org.jboss.seam.drools.ManagedWorkingMemory`

Manager component for a conversation scoped managed Drools `WorkingMemory`.

- `<managedWorkingMemory>.ruleBase` — a value expression that evaluates to an instance of `RuleBase`.

`<ruleBase>`, `org.jboss.seam.drools.RuleBase`

Manager component for an application scoped Drools `RuleBase`. *Note that this is not really intended for production usage, since it does not support dynamic installation of new rules.*

- `<ruleBase>.ruleFiles` — a list of files containing Drools rules.

`<ruleBase>.dslFile` — a Drools DSL definition.

`<entityHome>`, `org.jboss.seam.framework.EntityHome`

`<hibernateEntityHome>`, `org.jboss.seam.framework.HibernateEntityHome`

`<entityQuery>`, `org.jboss.seam.framework.EntityQuery`

`<hibernateEntityQuery>`, `org.jboss.seam.framework.HibernateEntityQuery`

Chapter 29. Seam JSF controls

Seam includes a number of JSF controls that are useful for working with Seam. These are intended to complement the built-in JSF controls, and controls from other third-party libraries. We recommend JBoss RichFaces, and Apache MyFaces Trinidad tag libraries for use with Seam. We do not recommend the use of the Tomahawk tag library.

29.1. Tags

To use these tags, define the "s" namespace in your page as follows (facelets only):

```
<html xmlns="http://www.w3.org/1999/xhtml"
      xmlns:s="http://jboss.com/products/seam/taglib">
```

The ui example demonstrates the use of a number of these tags.

Table 29.1. Seam JSF Control Reference

<code><s:button></code>	<i>Description</i> A button that supports invocation of an action with control over conversation propagation. <i>Does not submit the form.</i> <i>Attributes</i> • value — the label. • action — a method binding that specified the action listener. • view — the JSF view id to link to. • fragment — the fragment identifier to link to. • disabled — is the link disabled? • propagation — determines the conversation propagation style: begin, join, nest, none or end. • pageflow — a pageflow definition to begin. (This is only useful when propagation="begin" or propagation="join".) <i>Usage</i> <pre><s:button id="cancel" value="Cancel" action="#{hotelBooking.cancel}"/></pre>
<code><s:cache></code>	<i>Description</i> Cache the rendered page fragment using JBoss Cache. Note that <code><s:cache></code> actually uses the instance of JBoss Cache managed by the built-in <code>pojoCache</code> component.


	<i>Attributes</i> • key — the key to cache rendered content, often a value expression. For example, if we were caching a page fragment that displays a document, we might use <code>key="Document-#{document.id}"</code>. • enabled — a value expression that determines if the cache should be used. • region — a JBoss Cache node to use (different nodes can have different expiry policies). <i>Usage</i> <pre><s:cache key="entry-#{blogEntry.id}" region="pageFragments"> <div class="blogEntry"> <h3>#{blogEntry.title}</h3> <div> <s:formattedText value="#{blogEntry.body}" /> </div> <p> [Posted on#{160; <h:outputText value="#{blogEntry.date}"> <f:convertDateTime timezone="#{blog.timeZone}" locale="#{blog.locale}" type="both" /> </h:outputText>] </p> </div> </s:cache></pre>
<code><s:conversationId></code>	<i>Description</i> Add the conversation id to JSF link or button (e.g. <code><h:commandLink /></code>, <code><s:button /></code>). <i>Attributes</i> None.
<code><s:conversationPropagation></code>	<i>Description</i> Customize the conversation propagation for a command link or button (or similar JSF control). <i>Facelets only</i>. <i>Attributes</i> • type — determines the conversation propagation style: begin, join, nest, none or end. • pageflow — a pageflow definition to begin. (This is only useful when <code>propagation="begin"</code> or <code>propagation="join"</code>.) <i>Usage</i> <pre><h:commandButton value="Apply" action="#{personHome.update}"> <s:conversationPropagation type="join" /> </h:commandButton></pre>

<s:convertDateTime>	<i>Description</i> Perform date or time conversions in the Seam timezone. <i>Attributes</i> None. <i>Usage</i> <pre><h:outputText value="#{item.orderDate}"> <s:convertDateTime type="both" dateStyle="full"/> </h:outputText></pre>
<s:convertEntity>	<i>Description</i> Assigns an entity converter to the current component. This is primarily useful for radio button and dropdown controls. The converter works with any managed entity which has an @Id annotation - either simple or composite. <i>Attributes</i> None. <i>Configuration</i> You must use <i>Seam managed transactions</i> (see Section 8.2, “Seam managed transactions”) with <s:convertEntity />. If your <i>Managed Persistence Context</i> isn't called entityManager, then you need to set it in components.xml: <pre><component name="org.jboss.seam.ui.EntityConverter"> <property name="entityManager">#{em}</property> </component></pre> If you are using a <i>Managed Hibernate Session</i> then you need to set it in components.xml: <pre><component name="org.jboss.seam.ui.EntityConverter"> <property name="session">#{hibernateSession}</property> </component></pre> If you want to use more than one entity manager with the entity converter, you can create a copy of the entity converter for each entity manager; in components.xml: <pre><component name="myEntityConverter" class="org.jboss.seam.ui.converter.EntityConverter"> <property name="entityManager">#{em}</property> </component></pre> <pre><h:selectOneMenu value="#{person.continent}"> <s:selectItems value="#{continents.resultList}" var="continent" label="#{continent.name}" /> <f:converter converterId="myEntityConverter" /> </h:selectOneMenu></pre>

	<i>Usage</i> <pre><h:selectOneMenu value="#{person.continent}" required="true"> <s:selectItems value="#{continents.resultList}" var="continent" label="#{continent.name}" noSelectionLabel="Please Select..." /> <s:convertEntity /> </h:selectOneMenu></pre>
<code><s:convertEnum></code>	<i>Description</i> Assigns an enum converter to the current component. This is primarily useful for radio button and dropdown controls. <i>Attributes</i> None. <i>Usage</i> <pre><h:selectOneMenu value="#{person.honorific}"> <s:selectItems value="#{honorifics}" var="honorific" label="#{honorific.label}" noSelectionLabel="Please select" /> <s:convertEnum /> </h:selectOneMenu></pre>
<code><s:decorate></code>	<i>Description</i> "Decorate" a JSF input field when validation fails or when <code>required="true"</code> is set. <i>Attributes</i> <code>template</code> — the facelets template to use to decorate the component <code>#{invalid}</code> and <code>#{required}</code> are available inside <code>s:decorate</code>; <code>#{required}</code> evaluates to true if you have set the input component being decorated as required, and <code>#{invalid}</code> evaluates to true if a validation error occurs. <i>Usage</i> <pre><s:decorate template="edit.xhtml"> <ui:define name="label">Country:</ui:define> <h:inputText value="#{location.country}" required="true"/> </s:decorate></pre> <pre><ui:composition xmlns="http://www.w3.org/1999/xhtml" xmlns:ui="http://java.sun.com/jsf/facelets" xmlns:h="http://java.sun.com/jsf/html" xmlns:f="http://java.sun.com/jsf/core" xmlns:s="http://jboss.com/products/seam/taglib"> <div> <s:label styleClass="#{invalid?'error':''}"> <ui:insert name="label"/> <s:span styleClass="required" rendered="#{required}">*</s:span> </s:label></pre>

	<pre> <s:validateAll> <ui:insert/> </s:validateAll> <s:message styleClass="error"/> </div> </ui:composition> </pre>
<s:defaultAction>	<i>Description</i> Specify the default action to run when the form is submitted using the enter key. Currently you can only nest it inside buttons (e.g. <h:commandButton />, <a:commandButton /> or <tr:commandButton />). You must specify an id on the action source. You can only have one default action per form. <i>Attributes</i> None. <i>Usage</i> <pre> <h:commandButton id="foo" value="Foo" action="#{manager.foo}"> <s:defaultAction /> </h:commandButton> </pre>
<s:div>	<i>Description</i> Render a HTML <div>. <i>Attributes</i> None. <i>Usage</i> <pre> <s:div rendered="#{selectedMember == null}"> Sorry, but this member does not exist. </s:div> </pre>
<s:enumItem>	<i>Description</i> Creates a SelectItem from an enum value. <i>Attributes</i> - enumValue — the string representation of the enum value. - label — the label to be used when rendering the SelectItem. <i>Usage</i>

	<pre> <h:selectOneRadio id="radioList" layout="lineDirection" value="#{newPayment.paymentFrequency}"> <s:convertEnum /> <s:enumItem enumValue="ONCE" label="Only Once" /> <s:enumItem enumValue="EVERY_MINUTE" label="Every Minute" /> <s:enumItem enumValue="HOURLY" label="Every Hour" /> <s:enumItem enumValue="DAILY" label="Every Day" /> <s:enumItem enumValue="WEEKLY" label="Every Week" /> </h:selectOneRadio> </pre>
<s:fileUpload>	<i>Description</i> Renders a file upload control. This control must be used within a form with an encoding type of multipart/form-data, i.e: <pre> <h:form enctype="multipart/form-data"> </pre> For multipart requests, the Seam Multipart servlet filter must also be configured in web.xml: <pre> <filter> <filter-name>Seam Filter</filter-name> <filter-class>org.jboss.seam.servlet.SeamFilter</filter-class> </filter> <filter-mapping> <filter-name>Seam Filter</filter-name> <url-pattern>/*</url-pattern> </filter-mapping> </pre> <i>Configuration</i> The following configuration options for multipart requests may be configured in components.xml: • createTempFiles — if this option is set to true, uploaded files are streamed to a temporary file instead of in memory. • maxRequestSize — the maximum size of a file upload request, in bytes. Here's an example: <pre> <component class="org.jboss.seam.web.MultipartFilter"> <property name="createTempFiles">true</property> <property name="maxRequestSize">1000000</property> </component> </pre> <i>Attributes</i> • data — this value binding receives the binary file data. The receiving field should be declared as a byte[] or InputStream (required). • contentType — this value binding receives the file's content type (optional). • fileName — this value binding receives the filename (optional). • fileSize — this value binding receives the file size (optional).

	• <code>accept</code> — a comma-separated list of content types to accept, may not be supported by the browser. E.g. "images/png, images/jpg", "images/*". • <code>style</code> — The control's style • <code>styleClass</code> — The control's style class <i>Usage</i> <pre><s:fileUpload id="picture" data="#{register.picture}" accept="image/png" contentType="#{register.pictureContentType}" /></pre>
<code><s:formattedText></code>	<i>Description</i> Outputs <i>Seam Text</i>, a rich text markup useful for blogs, wikis and other applications that might use rich text. See the <i>Seam Text</i> chapter for full usage. <i>Attributes</i> • <code>value</code> — an EL expression specifying the rich text markup to render. <i>Usage</i> <pre><s:formattedText value="#{blog.text}" /></pre> <i>Example</i>  The example shows a web form titled "Please type your comment". The form contains a rich text editor with the following text: +Lorem ipsum *Lorem ipsum* /dolor sit amet/, consectetuer adipiscing elit . -Suspendisse a risus- ^quis lorem pharetra viverra^, _Fusce in ipsum. Nam et turpis id arcu lobortis dapibus_. ++Curabitur et sem vel quam #venenatis mattis. #Nulla hendrerit orci ut massa. The preview shows the rendered HTML: Preview Lorem ipsum Lorem ipsum dolor sit amet, consectetur adipiscing elit. -Suspendisse a risus- quis lorem pharetra viverra, Fusce in ipsum. Nam et turpis id arcu lobortis dapibus. Curabitur et sem vel quam 1. venenatis mattis. 2. Nulla hendrerit orci ut massa. 3. Donec condimentum, • libero in iaculis hendrerit, • risus dolor congue nulla, • non accumsan ante risus et ipsum. "Suspendisse dui. Maecenas lorem. Maecenas sit amet purus nec metus sodales sagittis. Phasellus varius lacus nec velit."
<code><s:fragment></code>	<i>Description</i> A non-rendering component useful for enabling/disabling rendering of it's children. <i>Attributes</i>

	None. <i>Usage</i> <pre><s:fragment rendered="#{auction.highBidder ne null}"> Current bid: </s:fragment></pre>
<s:graphicImage>	<i>Description</i> An extended <h:graphicImage> that allows the image to be created in a Seam Component; further transforms can be applied to the image. All attributes for <h:graphicImage> are supported, as well as: <i>Attributes</i> • <code>value</code> — image to display. Can be a path String (loaded from the classpath), a byte[], a <code>java.io.File</code>, a <code>java.io.InputStream</code> or a <code>java.net.URL</code>. Currently supported image formats are <code>image/png</code>, <code>image/jpeg</code> and <code>image/gif</code>. • <code>fileName</code> — if not specified the served image will have a generated file name. If you want to name your file, you should specify it here. This name should be unique <i>Transformations</i> To apply a transform to the image, you would nest a tag specifying the transform to apply. Seam currently supports these transforms: <s:transformImageSize> • <code>width</code> — new width of the image • <code>height</code> — new height of the image • <code>maintainRatio</code> — if <code>true</code>, and <i>one</i> of width/height are specified, the image will be resized with the dimension not specified being calculated to maintain the aspect ratio. • <code>factor</code> — scale the image by the given factor <s:transformImageBlur> • <code>radius</code> — perform a convolution blur with the given radius <s:transformImageType> • <code>contentType</code> — alter the type of the image to either <code>image/jpeg</code> or <code>image/png</code> It's easy to create your own transform - create a <code>UIComponent</code> which implements

	org.jboss.seam.ui.graphicImage.ImageTransform. Inside the applyTransform() method use image.getBufferedImage() to get the original image and image.setBufferedImage() to set your transformed image. Transforms are applied in the order specified in the view. <i>Usage</i> <pre><s:graphicImage rendered="#{auction.image ne null}" value="#{auction.image.data}"> <s:transformImageSize width="200" maintainRatio="true"/> </s:graphicImage></pre>
<s:label>	<i>Description</i> "Decorate" a JSF input field with the label. The label is placed inside the HTML <label> tag, and is associated with the nearest JSF input component. It is often used with <s:decorate>. <i>Attributes</i> - style — The control's style - styleClass — The control's style class <i>Usage</i> <pre><s:label styleClass="label"> Country: </s:label> <h:inputText value="#{location.country}" required="true"/></pre>
<s:link>	<i>Description</i> A link that supports invocation of an action with control over conversation propagation. <i>Does not submit the form.</i> <i>Attributes</i> - value — the label. - action — a method binding that specified the action listener. - view — the JSF view id to link to. - fragment — the fragment identifier to link to. - disabled — is the link disabled? - propagation — determines the conversation propagation style: begin, join, nest, none or end. - pageflow — a pageflow definition to begin. (This is only useful when propagation="begin" or propagation="join".) <i>Usage</i>

	<pre><s:link id="register" view="/register.xhtml" value="Register New User"/></pre>
<s:message>	<i>Description</i> "Decorate" a JSF input field with the validation error message. <i>Attributes</i> None. <i>Usage</i> <pre><f:facet name="afterInvalidField"> <s:span> &#160;Error:&#160; <s:message/> </s:span> </f:facet></pre>
<s:remote>	<i>Description</i> Generates the Javascript stubs required to use Seam Remoting. <i>Attributes</i> include — a comma-separated list of the component names (or fully qualified class names) for which to generate Seam Remoting Javascript stubs. See Chapter 21, <i>Remoting</i> for more details. <i>Usage</i> <pre><s:remote include="customerAction,accountAction,com.acme.MyBean" /></pre>
<s:selectDate>	<i>Description</i> <i>Deprecated. Use <rich:calendar /> instead.</i> Displays a dynamic date picker component that selects a date for the specified input field. The body of the selectDate element should contain HTML elements, such as text or an image, that prompt the user to click to display the date picker. The date picker <i>must</i> be styled using CSS. An example CSS file can be found in the Seam booking demo as date.css, or can be generated using seam-gen. The CSS styles used to control the appearance of the date picker are also described below. <i>Attributes</i> - for — The id of the input field that the date picker will insert the selected date into. - dateFormat — The date format string. This should match the date format of the input field.

- `startYear` — The popup year selector range will start at this year.
- `endYear` — The popup year selector range will end at this year.
- `firstDayOfWeek` — Controls which day is the first day of the week (0 = Sunday, 6 = Saturday). If this attribute is not set, then the first day of the week will default based on the user's locale.

Usage

```
<div class="row">
  <h:outputLabel for="dob">Date of birth<em>*</em></h:outputLabel>
  <h:inputText id="dob" value="#{user.dob}" required="true">
    <s:convertDateTime pattern="MM/dd/yyyy"/>
  </h:inputText>
  <s:selectDate for="dob" startYear="1910" endYear="2007">
    
  </s:selectDate>
  <div class="validationError"><h:message for="dob"/></div>
</div>
```

Example




CSS Styling

The following list describes the CSS class names that are used to control the style of the `selectDate` control.

- `seam-date` — This class is applied to the outer `div` containing the popup calendar. (1) It is also applied to the `table` that controls the inner layout of the calendar. (2)
- `seam-date-header` — This class is applied to the calendar header table row (`tr`) and header table cells (`td`). (3)
- `seam-date-header-prevMonth` — This class is applied to the "previous month" table cell, (`td`), which when clicked causes the calendar to display the month prior to the one currently displayed. (4)

- `seam-date-header-nextMonth` — This class is applied to the "next month" table cell, (td), which when clicked causes the calendar to display the month following the one currently displayed. (5)
- `seam-date-headerDays` — This class is applied to the calendar days header row (tr), which contains the names of the week days. (6)
- `seam-date-footer` — This class is applied to the calendar footer row (tr), which displays the current date. (7)
- `seam-date-inMonth` — This class is applied to the table cell (td) elements that contain a date within the month currently displayed. (8)
- `seam-date-outMonth` — This class is applied to the table cell (td) elements that contain a date outside of the month currently displayed. (9)
- `seam-date-selected` — This class is applied to the table cell (td) element that contains the currently selected date. (10)
- `seam-date-dayOff-inMonth` — This class is applied to the table cell (td) elements that contain a "day off" date (i.e. weekend days, Saturday and Sunday) within the currently selected month. (11)
- `seam-date-dayOff-outMonth` — This class is applied to the table cell (td) elements that contain a "day off" date (i.e. weekend days, Saturday and Sunday) outside of the currently selected month. (12)
- `seam-date-hover` — This class is applied to the table cell (td) element over which the cursor is hovering. (13)
- `seam-date-monthNames` — This class is applied to the div control that contains the popup month selector. (14)
- `seam-date-monthNameLink` — This class is applied to the anchor (a) controls that contain the popup month names. (15)
- `seam-date-years` — This class is applied to the div control that contains the popup year selector. (16)
- `seam-date-yearLink` — This class is applied to the anchor (a) controls that contain the popup years. (17)



<s:selectItems>

Description

Creates a `List<SelectItem>` from a `List`, `Set`, `DataModel` or `Array`.

	<i>Attributes</i> • <code>value</code> — an EL expression specifying the data that backs the <code>List<SelectItem></code> • <code>var</code> — defines the name of the local variable that holds the current object during iteration • <code>label</code> — the label to be used when rendering the <code>SelectItem</code>. Can reference the <code>var</code> variable • <code>itemValue</code> — Value to return to the server if this option is selected. Optional, by default the <code>var</code> object is used. Can reference the <code>var</code> variable • <code>disabled</code> — if true the <code>SelectItem</code> will be rendered disabled. Can reference the <code>var</code> variable • <code>noSelectionLabel</code> — specifies the (optional) label to place at the top of list (if <code>required="true"</code> is also specified then selecting this value will cause a validation error) • <code>hideNoSelectionLabel</code> — if true, the <code>noSelectionLabel</code> will be hidden when a value is selected <i>Usage</i> <pre><h:selectOneMenu value="#{person.age}" converter="#{converters.ageConverter}"> <s:selectItems value="#{ages}" var="age" label="#{age}" /> </h:selectOneMenu></pre>
<code><s:span></code>	<i>Description</i> Render a HTML <code></code>. <i>Attributes</i> None. <i>Usage</i> <pre><s:span styleClass="required" rendered="#{required}">*</s:span></pre>
<code><s:taskId></code>	<i>Description</i> Add the task id to an output link (or similar JSF control), when the task is available via <code>{task}</code>. <i>Attributes</i> None.
<code><s:validate></code>	<i>Description</i>

	A non-visual control, validates a JSF input field against the bound property using Hibernate Validator. <i>Attributes</i> None. <i>Usage</i> <pre><h:inputText id="userName" required="true" value="#{customer.userName}"> <s:validate /> </h:inputText> <h:message for="userName" styleClass="error" /></pre>
<s:validateAll>	<i>Description</i> A non-visual control, validates all child JSF input fields against their bound properties using Hibernate Validator. <i>Attributes</i> None. <i>Usage</i> <pre><s:validateAll> <div class="entry"> <h:outputLabel for="username">Username:</h:outputLabel> <h:inputText id="username" value="#{user.username}" required="true"/> <h:message for="username" styleClass="error" /> </div> <div class="entry"> <h:outputLabel for="password">Password:</h:outputLabel> <h:inputSecret id="password" value="#{user.password}" required="true"/> <h:message for="password" styleClass="error" /> </div> <div class="entry"> <h:outputLabel for="verify">Verify Password:</h:outputLabel> <h:inputSecret id="verify" value="#{register.verify}" required="true"/> <h:message for="verify" styleClass="error" /> </div> </s:validateAll></pre>
<s:validateFormattedText>	<i>Description</i> Checks that the submitted value is valid Seam Text <i>Attributes</i> None.

29.2. Annotations

Seam also provides annotations to allow you to use Seam components as JSF converters and validators:

@Converter

```
@Name("fooConverter")
@BypassInterceptors
@Converter
public class FooConverter implements Converter {

    @Transactional
    public Object getAsObject(FacesContext context, UIComponent cmp, String value) {
        EntityManager entityManager = (EntityManager) Component.getInstance("entityManager");
        entityManager.joinTransaction();
        // Do the conversion
    }

    public String getAsString(FacesContext context, UIComponent cmp, Object value) {
        // Do the conversion
    }
}
```

Registers the Seam component as a JSF converter. Shown here is a converter which is able to access the JPA EntityManager inside a JTA transaction, when converting the value back to it's object representation.

@Validator

```
@Name("barValidator")
@BypassInterceptors
@Validator
public class BarValidator implements Validator {

    public void validate(FacesContext context, UIComponent cmp, Object value)
        throws ValidatorException {
        FooController fooController = (FooController) Component.getInstance("fooController");
        return fooController.validate(value);
    }
}
```

Registers the Seam component as a JSF validator. Shown here is a validator which injects another Seam component; the injected component is used to validate the value.

Chapter 30. Expression language enhancements

Seam provides an extension to the standard Unified Expression Language (EL) called JBoss EL. JBoss EL provides a number of enhancements that increase the expressiveness and power of EL expressions.

30.1. Parameterized Method Bindings

Standard EL assumes that any parameters to a method expression will be provided by Java code. This means that a method with parameters cannot be used as a JSF method binding. Seam provides an enhancement to the EL that allows parameters to be included in a method expression itself. This applies to *any* Seam method expression, including any JSF method binding, for example:

```
<h:commandButton action="#{hotelBooking.bookHotel(hotel)}" value="Book Hotel"/>
```

30.1.1. Usage

Parameters are surrounded by parentheses, and separated by commas:

```
<h:commandButton action="#{hotelBooking.bookHotel(hotel, user)}" value="Book Hotel"/>
```

The parameters `hotel` and `user` will be evaluated as value expressions and passed to the `bookHotel()` method of the component. This gives you an alternative to the use of `@In`.

Any value expression may be used as a parameter:

```
<h:commandButton action="#{hotelBooking.bookHotel(hotel.id, user.username)}"
                 value="Book Hotel"/>
```

Note: You can not pass objects as arguments! All that is passed is names, for example, `hotel.id` and `user.username`. If you check the rendered code of the previous example, you will see that the command button contains these names. These name arguments will be submitted to the server when you press the button, and Seam will look up and resolve these names (in any available context) before the action method is called. If the arguments can not be resolved at that time (because `hotel` and `user` variables can not be found in any available context) the action method will be called with `null` arguments!

You may however pass literal strings using single or double quotes:

```
<h:commandLink action="#{printer.println('Hello world!')}" value="Hello"/>
```

```
<h:commandLink action="#{printer.println('Hello again')}" value="Hello"/>
```

You might even want to use this notation for all your action methods, even when you don't have parameters to pass. This improves readability by making it clear that the expression is a method expression and not a value expression:

```
<s:link value="Cancel" action="#{hotelBooking.cancel()}" />
```

30.1.2. Limitations

Please be aware of the following limitations:

30.1.2.1. Incompatibility with JSP 2.1

This extension is not currently compatible with JSP 2.1. So if you want to use this extension with JSF 1.2, you will need to use Facelets. The extension works correctly with JSP 2.0.

30.1.2.2. Calling a `MethodExpression` from Java code

Normally, when a `MethodExpression` or `MethodBinding` is created, the parameter types are passed in by JSF. In the case of a method binding, JSF assumes that there are no parameters to pass. With this extension, we can't know the parameter types until after the expression has been evaluated. This has two minor consequences:

- When you invoke a `MethodExpression` in Java code, parameters you pass may be ignored. Parameters defined in the expression will take precedence.
- Ordinarily, it is safe to call `methodExpression.getMethodInfo().getParamTypes()` at any time. For an expression with parameters, you must first invoke the `MethodExpression` before calling `getParamTypes()`.

Both of these cases are exceedingly rare and only apply when you want to invoke the `MethodExpression` by hand in Java code.

30.2. Parameterized Value Bindings

Standard EL only allows access to properties that follow the JavaBean naming conventions. For example, the expression `#{person.name}` requires a `getName()` be present. Many objects, however, don't have appropriately named property accessors or require parameters. These values can be retrieved using the method syntax, which work similarly to parameterized method bindings. For example, the following expression returns the size of a string using the `length()` method.

```
#{person.name.length()}
```

You can access the size of a collection in a similar manner.

```
#{searchResults.size()}
```

In general any expression of the form `#{obj.property}` would be identical to the expression `#{obj.getProperty()}`.

Parameters are also allowed, and they follow the same restrictions as with method bindings. The following example calls the `productsByColorMethod` with a literal string argument.

```
#{controller.productsByColor('blue')}
```

30.3. Projection

JBoss EL supports a limited projection syntax. It is important to note that this syntax cannot be parsed by Facelets or by JavaServer Pages and thus cannot be used in xhtml or JSP files. We anticipate that the projection syntax will change in future versions of JBoss EL.

A projection expression maps a sub-expression across a multi-valued (list, set, etc...) expression. For instance, the expression

```
#{company.departments}
```

might return a list of departments. If you only need a list of department names, your only option is to iterate over the list to retrieve the values. JBoss EL allows this with a projection expression.

```
#{company.departments.{d|d.name}}
```

The subexpression is enclosed in braces. In this example, the expression `d.name` is evaluated for each department, using `d` as an alias to the department object. The result of this expression will be a list of String values.

Any valid expression can be used in an expression, so it would be perfectly valid to write the following, assuming you had a use for the lengths of all the department names in a company.

```
#{company.departments.{d|d.size()}}
```

Projections can be nested. The following expression returns the last names of every employee in every department.

```
#{company.departments.{d|d.employees.{emp|emp.lastName}}}
```

Nested projections can be slightly tricky, however. The following expression looks like it returns a list of all the employees in all the departments.

```
#{company.departments.{d|d.employees}}
```

However, it actually returns a list containing a list of the employees for each individual department. To combine the values, it is necessary to use a slightly longer expression.

```
#{company.departments.{d|d.employees.{e|e}}}
```

Chapter 31. Testing Seam applications

Most Seam applications will need at least two kinds of automated tests: *unit tests*, which test a particular Seam component in isolation, and scripted *integration tests* which exercise all Java layers of the application (that is, everything except the view pages).

Both kinds of tests are very easy to write.

31.1. Unit testing Seam components

All Seam components are POJOs. This is a great place to start if you want easy unit testing. And since Seam emphasises the use of bijection for inter-component interactions and access to contextual objects, it's very easy to test a Seam component outside of its normal runtime environment.

Consider the following Seam component:

```
@Stateless
@Scope(EVENT)
@Name("register")
public class RegisterAction implements Register
{
    private User user;
    private EntityManager em;

    @In
    public void setUser(User user) {
        this.user = user;
    }

    @PersistenceContext
    public void setBookingDatabase(EntityManager em) {
        this.em = em;
    }

    public String register()
    {
        List existing = em.createQuery("select username from User where username=:username")
            .setParameter("username", user.getUsername())
            .getResultList();
        if (existing.size()==0)
        {
            em.persist(user);
            return "success";
        }
        else
        {
            return null;
        }
    }
}
```

We could write a TestNG test for this component as follows:

```
public class RegisterActionTest
{
    @Test
    public void testRegisterAction()
    {
        EntityManager em = getEntityManagerFactory().createEntityManager();
        em.getTransaction().begin();
    }
}
```

```

        User gavin = new User();
        gavin.setName("Gavin King");
        gavin.setUserName("lovthafew");
        gavin.setPassword("secret");

        RegisterAction action = new RegisterAction();
        action.setUser(gavin);
        action.setBookingDatabase(em);

        assert "success".equals( action.register() );

        em.getTransaction().commit();
        em.close();
    }

    private EntityManagerFactory emf;

    public EntityManagerFactory getEntityManagerFactory()
    {
        return emf;
    }

    @BeforeClass
    public void init()
    {
        emf = Persistence.createEntityManagerFactory("myResourceLocalEntityManager");
    }

    @AfterClass
    public void destroy()
    {
        emf.close();
    }
}

```

The Java Persistence API can be used with both Java SE and Java EE — when the above component is used inside an Application Server (Java EE) the container is responsible for transaction management, however in the unit test (Java SE) the transaction must be managed explicitly using a resource local entity manager. This requires configuration in `persistence.xml`.

Seam components don't usually depend directly upon container infrastructure, so most unit testing is as easy as that!

31.2. Integration testing Seam components

Integration testing is slightly more difficult. In this case, we can't eliminate the container infrastructure; indeed, that is part of what is being tested! At the same time, we don't want to be forced to deploy our application to an application server to run the automated tests. We need to be able to reproduce just enough of the container infrastructure inside our testing environment to be able to exercise the whole application, without hurting performance too much.

The approach taken by Seam is to let you write tests that exercise your components while running inside a pruned down container environment (Seam, together with the JBoss Embedded container; n.b. JBoss Embedded requires JDK 1.5 and does not work with JDK 1.6).

```

public class RegisterTest extends SeamTest
{
    @Test

```

```

public void testRegisterComponent() throws Exception
{
    new ComponentTest() {
        protected void testComponents() throws Exception
        {
            setValue("#{user.username}", "1ovthafew");
            setValue("#{user.name}", "Gavin King");
            setValue("#{user.password}", "secret");
            assert invokeMethod("#{register.register}").equals("success");
            assert getValue("#{user.username}").equals("1ovthafew");
            assert getValue("#{user.name}").equals("Gavin King");
            assert getValue("#{user.password}").equals("secret");
        }
    }.run();
}
...
}

```

31.2.1. Using mocks in integration tests

Occasionally, we need to be able to replace the implementation of some Seam component that depends upon resources which are not available in the integration test environment. For example, suppose we have some Seam component which is a facade to some payment processing system:

```

@Name("paymentProcessor")
public class PaymentProcessor {
    public boolean processPayment(Payment payment) { .... }
}

```

For integration tests, we can mock out this component as follows:

```

@Name("paymentProcessor")
@Install(precedence=MOCK)
public class MockPaymentProcessor extends PaymentProcessor {
    public void processPayment(Payment payment) {
        return true;
    }
}

```

Since the `MOCK` precedence is higher than the default precedence of application components, Seam will install the mock implementation whenever it is in the classpath. When deployed into production, the mock implementation is absent, so the real component will be installed.

31.3. Integration testing Seam application user interactions

An even harder problem is emulating user interactions. A third problem is where to put our assertions. Some test frameworks let us test the whole application by reproducing user interactions with the web browser. These frameworks have their place, but they are not appropriate for use at development time.

SeamTest lets you write *scripted* tests, in a simulated JSF environment. The role of a scripted test is to reproduce the interaction between the view and the Seam components. In other words, you get to pretend you are the JSF implementation!

This approach tests everything except the view.

Let's consider a JSP view for the component we unit tested above:

```
<html>
<head>
  <title>Register New User</title>
</head>
<body>
  <f:view>
    <h:form>
      <table border="0">
        <tr>
          <td>Username</td>
          <td><h:inputText value="#{user.username}"/></td>
        </tr>
        <tr>
          <td>Real Name</td>
          <td><h:inputText value="#{user.name}"/></td>
        </tr>
        <tr>
          <td>Password</td>
          <td><h:inputSecret value="#{user.password}"/></td>
        </tr>
      </table>
      <h:messages/>
      <h:commandButton type="submit" value="Register" action="#{register.register}"/>
    </h:form>
  </f:view>
</body>
</html>
```

We want to test the registration functionality of our application (the stuff that happens when the user clicks the Register button). We'll reproduce the JSF request lifecycle in an automated TestNG test:

```
public class RegisterTest extends SeamTest
{
    @Test
    public void testRegister() throws Exception
    {
        new FacesRequest() {
            @Override
            protected void processValidations() throws Exception
            {
                validateValue("#{user.username}", "1ovthafew");
                validateValue("#{user.name}", "Gavin King");
                validateValue("#{user.password}", "secret");
                assert !isValidationFailure();
            }

            @Override
            protected void updateModelValues() throws Exception
            {
                setValue("#{user.username}", "1ovthafew");
                setValue("#{user.name}", "Gavin King");
                setValue("#{user.password}", "secret");
            }

            @Override
            protected void invokeApplication()
            {
                assert invokeMethod("#{register.register}").equals("success");
            }

            @Override
            protected void renderResponse()
            {
            }
        };
    }
}
```


```
        {
            assert getValue("#{user.username}").equals("1ovthafew");
            assert getValue("#{user.name}").equals("Gavin King");
            assert getValue("#{user.password}").equals("secret");
        }

        }.run();
    }

    ...
}
```

Notice that we've extended `SeamTest`, which provides a Seam environment for our components, and written our test script as an anonymous class that extends `SeamTest.FacesRequest`, which provides an emulated JSF request lifecycle. (There is also a `SeamTest.NonFacesRequest` for testing GET requests.) We've written our code in methods which are named for the various JSF phases, to emulate the calls that JSF would make to our components. Then we've thrown in various assertions.

You'll find plenty of integration tests for the Seam example applications which demonstrate more complex cases. There are instructions for running these tests using Ant, or using the TestNG plugin for eclipse:



31.3.1. Integration Testing with Mock Data

If you need to insert or clean data in your database before each test you can use Seam's integration with DBUnit. To do this, extend DBUnitSeamTest rather than SeamTest.

You need to provide a dataset for DBUnit:

```
<dataset>

  <ARTIST
    id="1"
    dtype="Band"
    name="Pink Floyd" />

  <DISC
    id="1"
    name="Dark Side of the Moon"
    artist_id="1" />

</dataset>
```

and tell Seam about it by overriding `prepareDBUnitOperations()`:

```
protected void prepareDBUnitOperations() {
    beforeTestOperations.add(
        new DataSetOperation("my/datasets/BaseData.xml")
    );
}
```

`DataSetOperation` defaults to `DatabaseOperation.CLEAN_INSERT` if no other operation is specified as a constructor argument. The above example cleans all tables defined `BaseData.xml`, then inserts all rows declared in `BaseData.xml` before each `@Test` method is invoked.

If you require extra cleanup after a test method executes, add operations to `afterTestOperations` list.

You need to tell DBUnit about the datasource you are using by setting a TestNG test parameter named `datasourceJndiName`:

```
<parameter name="datasourceJndiName" value="java:/seamdiscsDatasource"/>
```

31.3.2. Integration Testing Seam Mail

Warning! This feature is still under development.

It's very easy to integration test your Seam Mail:

```
public class MailTest extends SeamTest {

    @Test
    public void testSimpleMessage() throws Exception {

        new FacesRequest() {

            @Override
            protected void updateModelValues() throws Exception {
                setValue("#{person.firstname}", "Pete");
                setValue("#{person.lastname}", "Muir");
                setValue("#{person.address}", "test@example.com");
            }

            @Override
            protected void invokeApplication() throws Exception {
```

```
        MimeMessage renderedMessage = getRenderedMailMessage("/simple.xhtml");
        assert renderedMessage.getAllRecipients().length == 1;
        InternetAddress to = (InternetAddress) renderedMessage.getAllRecipients()[0];
        assert to.getAddress().equals("test@example.com");
    }
    }.run();
}
```

We create a new `FacesRequest` as normal. Inside the `invokeApplication` hook we render the message using `getRenderedMailMessage(viewId)`, passing the `viewId` of the message to render. The method returns the rendered message on which you can do your tests. You can of course also use any of the standard JSF lifecycle methods.

There is no support for rendering standard JSF components so you can't test the content body of the mail message easily.


Chapter 32. Seam tools

32.1. jBPM designer and viewer

The jBPM designer and viewer will let you design and view in a nice way your business processes and your pageflows. This convenient tool is part of JBoss Eclipse IDE and more details can be found in the jBPM's documentation (<http://docs.jboss.com/jbpm/v3/gpd/>)


32.1.1. Business process designer

This tool lets you design your own business process in a graphical way.



32.1.2. Pageflow viewer

This tool let you design to some extend your pageflows and let you build graphical views of them so you can easily share and compare ideas on how it should be designed.



Chapter 33. Dependencies

33.1. Project Dependencies

This section both lists the compile-time and runtime dependencies for Seam. Where the type is listed as `ear`, the library should be included in the `/lib` directory of your application's ear file. Where the type is listed as `war`, the library should be placed in the `/WEB-INF/lib` directory of your application's war file. The scope of the dependency is either `all`, `runtime` or `provided` (by JBoss AS 4.2).

Up to date version information is not included in the docs, but is provided in the `/build/root.pom.xml` Maven POM.

33.1.1. Core

Table 33.1.

Name	Scope	Type	Notes
<code>commons-codec.jar</code>	<code>runtime</code>	<code>ear</code>	Required by Seam Security when using Digest authentication.
<code>jboss-seam.jar</code>	<code>all</code>	<code>ear</code>	The core Seam library, always required.
<code>jboss-seam-debug.jar</code>	<code>runtime</code>	<code>war</code>	Include during development when enabling Seam's debug feature
<code>jboss-seam-ioc.jar</code>	<code>runtime</code>	<code>war</code>	Required when using Seam with Spring
<code>jboss-seam-pdf.jar</code>	<code>runtime</code>	<code>war</code>	Required when using Seam's PDF features
<code>jboss-seam-remoting.jar</code>	<code>runtime</code>	<code>war</code>	Required when using Seam Remoting
<code>jboss-seam-ui.jar</code>	<code>runtime</code>	<code>war</code>	Required to use the Seam JSF controls
<code>jsf-api.jar</code>	<code>provided</code>		JSF API
<code>jsf-impl.jar</code>	<code>provided</code>		JSF Reference Implementation
<code>jsf-facelets.jar</code>	<code>runtime</code>	<code>war</code>	Facelets
<code>urlrewrite.jar</code>	<code>runtime</code>	<code>war</code>	URL Rewrite library

Name	Scope	Type	Notes
jcaptcha-all.jar	runtime	ear	Required for Captcha support
quartz.jar	runtime	ear	Required when you wish to use Quartz with Seam's asynchronous features

33.1.2. RichFaces

Table 33.2. RichFaces dependencies

Name	Scope	Type	Notes
richfaces-api.jar	all	ear	Required to use RichFaces. Provides API classes that you may wish to use from your application e.g. to create a tree
richfaces-impl.jar	runtime	war	Required to use RichFaces.
richfaces-ui.jar	runtime	war	Required to use RichFaces. Provides all the UI components.

33.1.3. Seam Mail

Table 33.3. Seam Mail Dependencies

Name	Scope	Type	Notes
activation.jar	runtime	ear	Required for attachment support
mail.jar	runtime	ear	Required for outgoing mail support
mail-ra.jar	compile only		Required for incoming mail support mail-ra.rar should be deployed to the application server at runtime
jboss-seam-mail.jar	runtime	war	Seam Mail

33.1.4. Seam PDF

Table 33.4. Seam PDF Dependencies

Name	Type	Scope	Notes
itext.jar	runtime	war	PDF Library
jfreechart.jar	runtime	war	Charting library
jcommon.jar	runtime	war	Required by JFreeChart
jboss-seam-pdf.jar	runtime	war	Seam PDF core library

33.1.5. JBoss Rules

The JBoss Rules libraries can be found in the drools/lib directory in Seam.

Table 33.5. JBoss Rules Dependencies

Name	Scope	Type	Notes
antlr-runtime.jar	runtime	ear	ANTLR Runtime Library
core.jar	runtime	ear	Eclipse JDT
drools-compiler.jar	runtime	ear	
drools-core.jar	runtime	ear	
janino.jar	runtime	ear	
mvel.jar	runtime	ear	

33.1.6. JBPM

Table 33.6. JBPM dependencies

Name	Scope	Type	Notes
jbpm-jpd1.jar	runtime	ear	

33.1.7. GWT

These libraries are required if you wish to use the Google Web Toolkit (GWT) with your Seam application.

Table 33.7. GWT dependencies

Name	Scope	Type	Notes
gwt-servlet.jar	runtime	war	The GWT Servlet libs

33.1.8. Spring

These libraries are required if you wish to use the Spring Framework with your Seam application.

Table 33.8. Spring Framework dependencies

Name	Scope	Type	Notes
spring.jar	runtime	ear	The Spring Framework library

33.1.9. Groovy

These libraries are required if you wish to use Groovy with your Seam application.

Table 33.9. Groovy dependencies

Name	Scope	Type	Notes
groovy-all.jar	runtime	ear	The Groovy libs

33.2. Dependency Management using Maven

Maven offers support for transitive dependency management and can be used to manage the dependencies of your Seam project. You can use Maven Ant Tasks to integrate Maven into your Ant build, or can use Maven to build and deploy your project.

We aren't actually going to discuss how to use Maven here, but just run over some basic POMs you could use.

Released versions of Seam are available in <http://repository.jboss.org/maven2> and nightly snapshots are available in <http://snapshots.jboss.org/maven2>.

All the Seam artifacts are available in Maven:

```
<dependency>
```

```
<groupId>org.jboss.seam</groupId>
<artifactId>jboss-seam</artifactId>
</dependency>
```

```
<dependency>
  <groupId>org.jboss.seam</groupId>
  <artifactId>jboss-seam-ui</artifactId>
</dependency>
```

```
<dependency>
  <groupId>org.jboss.seam</groupId>
  <artifactId>jboss-seam-pdf</artifactId>
</dependency>
```

```
<dependency>
  <groupId>org.jboss.seam</groupId>
  <artifactId>jboss-seam-remoting</artifactId>
</dependency>
```

```
<dependency>
  <groupId>org.jboss.seam</groupId>
  <artifactId>jboss-seam-ioc</artifactId>
</dependency>
```

```
<dependency>
  <groupId>org.jboss.seam</groupId>
  <artifactId>jboss-seam-ioc</artifactId>
</dependency>
```

This sample POM will give you Seam, JPA (provided by Hibernate) and Hibernate Validator:

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
  <modelVersion>4.0.0</modelVersion>
  <groupId>org.jboss.seam.example</groupId>
  <artifactId>my-project</artifactId>
  <version>1.0</version>
  <name>My Seam Project</name>
  <packaging>jar</packaging>
  <repositories>
    <repository>
      <id>repository.jboss.org</id>
      <name>JBoss Repository</name>
      <url>http://repository.jboss.org/maven2</url>
    </repository>
  </repositories>

  <dependencies>

    <dependency>
      <groupId>org.hibernate</groupId>
      <artifactId>hibernate-validator</artifactId>
      <version>3.0.0.GA</version>
    </dependency>

    <dependency>
      <groupId>org.hibernate</groupId>
      <artifactId>hibernate-annotations</artifactId>
      <version>3.3.0.ga</version>
    </dependency>

    <dependency>
      <groupId>org.hibernate</groupId>
      <artifactId>hibernate-entitymanager</artifactId>
```

```
    <version>3.3.1.ga</version>
  </dependency>

  <dependency>
    <groupId>org.jboss.seam</groupId>
    <artifactId>jboss-seam</artifactId>
    <version>2.0.0.GA</version>
  </dependency>

</dependencies>

</project>
```