

**Web Beans: Java Contexts
and Dependency Injection**

**The new standard
for dependency
injection and contextual
state management**

Gavin King

**JSR-299 specification lead
Red Hat Middleware LLC**

Pete Muir

**Web Beans (JSR-299 Reference Implementation) lead
Red Hat Middleware LLC**

Italian Translation: Nicola Benaglia, Francesco Milesi

**Spanish Translation: Gladys Guerrero
Red Hat Middleware LLC**

**Korean Translation: Eun-Ju Ki,
Red Hat Middleware LLC**

Traditional Chinese Translation: Terry Chuang

Red Hat Middleware LLC

Note	vii
I. Using contextual objects	1
1. Getting started with Web Beans	3
1.1. Your first Web Bean	3
1.2. What is a Web Bean?	5
1.2.1. API types, binding types and dependency injection	6
1.2.2. Deployment types	7
1.2.3. Scope	8
1.2.4. Web Bean names and Unified EL	8
1.2.5. Interceptor binding types	9
1.3. What kinds of objects can be Web Beans?	9
1.3.1. Simple Web Beans	9
1.3.2. Enterprise Web Beans	10
1.3.3. Producer methods	11
1.3.4. JMS endpoints	12
2. JSF web application example	13
3. The Web Beans Reference Implementation	17
3.1. The numberguess example	19
3.2. The translator example	26
4. Dependency injection	31
4.1. Binding annotations	33
4.1.1. Binding annotations with members	34
4.1.2. Combinations of binding annotations	34
4.1.3. Binding annotations and producer methods	35
4.1.4. The default binding type	35
4.2. Deployment types	35
4.2.1. Enabling deployment types	36
4.2.2. Deployment type precedence	37
4.2.3. Example deployment types	37
4.3. Fixing unsatisfied dependencies	38
4.4. Client proxies	38
4.5. Obtaining a Web Bean by programatic lookup	39
4.6. Lifecycle callbacks, <code>@Resource</code> , <code>@EJB</code> and <code>@PersistenceContext</code>	40
4.7. The <code>InjectionPoint</code> object	40
5. Scopes and contexts	43
5.1. Scope types	43
5.2. Built-in scopes	43
5.3. The conversation scope	44
5.3.1. Conversation demarcation	45
5.3.2. Conversation propagation	46
5.3.3. Conversation timeout	46
5.4. The dependent pseudo-scope	46
5.4.1. The <code>@New</code> annotation	47
6. Producer methods	49

6.1. Scope of a producer method	50
6.2. Injection into producer methods	50
6.3. Use of @New with producer methods	51
II. Developing loosely-coupled code	53
7. Interceptors	55
7.1. Interceptor bindings	55
7.2. Implementing interceptors	56
7.3. Enabling interceptors	56
7.4. Interceptor bindings with members	57
7.5. Multiple interceptor binding annotations	58
7.6. Interceptor binding type inheritance	59
7.7. Use of @Interceptors	60
8. Decorators	61
8.1. Delegate attributes	62
8.2. Enabling decorators	62
9. Events	65
9.1. Event observers	65
9.2. Event producers	66
9.3. Registering observers dynamically	67
9.4. Event bindings with members	67
9.5. Multiple event bindings	68
9.6. Transactional observers	69
III. Making the most of strong typing	73
10. Stereotypes	75
10.1. Default scope and deployment type for a stereotype	75
10.2. Restricting scope and type with a stereotype	76
10.3. Interceptor bindings for stereotypes	77
10.4. Name defaulting with stereotypes	77
10.5. Standard stereotypes	78
11. Specialization	79
11.1. Using specialization	80
11.2. Advantages of specialization	80
12. Defining Web Beans using XML	83
12.1. Declaring Web Bean classes	83
12.2. Declaring Web Bean metadata	84
12.3. Declaring Web Bean members	85
12.4. Declaring inline Web Beans	85
12.5. Using a schema	86
IV. Web Beans and the Java EE ecosystem	87
13. Java EE integration	89
13.1. Injecting Java EE resources into a Web Bean	89
13.2. Calling a Web Bean from a Servlet	90
13.3. Calling a Web Bean from a Message-Driven Bean	90
13.4. JMS endpoints	91

13.5. Packaging and deployment	92
14. Extending Web Beans	93
14.1. The <code>Manager</code> object	93
14.2. The <code>Bean</code> class	94
14.3. The <code>Context</code> interface	95
15. Next steps	97
A. Integrating the Web Beans RI into other environments	99
A.1. The Web Beans RI SPI	99
A.1.1. Web Bean Discovery	99
A.1.2. EJB Discovery	100
A.1.3. JNDI	103
A.1.4. Resource loading	104
A.2. The contract with the container	105

Note

JSR-299 has recently changed its name from "Web Beans" to "Java Contexts and Dependency Injection". The reference guide still refers to JSR-299 as "Web Beans" and the JSR-299 Reference Implementation as the "Web Beans RI". Other documentation, blogs, forum posts etc. may use the new nomenclature, including the new name for the JSR-299 Reference Implementation - "Web Beans".

Part I. Using contextual objects

The Web Beans (JSR-299) specification defines a set of services for the Java EE environment that makes applications much easier to develop. Web Beans layers an enhanced lifecycle and interaction model over existing Java component types including JavaBeans and Enterprise Java Beans. As a complement to the traditional Java EE programming model, the Web Beans services provide:

- an improved lifecycle for stateful components, bound to well-defined *contexts*,
- a typesafe approach to *dependency injection*,
- interaction via an *event notification* facility, and
- a better approach to binding *interceptors* to components, along with a new kind of interceptor, called a *decorator*, that is more appropriate for use in solving business problems.

Dependency injection, together with contextual lifecycle management, saves the user of an unfamiliar API from having to ask and answer the following questions:

- what is the lifecycle of this object?
- how many simultaneous clients can it have?
- is it multithreaded?
- where can I get one from?
- do I need to explicitly destroy it?
- where should I keep my reference to it when I'm not using it directly?
- how can I add an indirection layer, so that the implementation of this object can vary at deployment time?
- how should I go about sharing this object between other objects?

A Web Bean specifies only the type and semantics of other Web Beans it depends upon. It need not be aware of the actual lifecycle, concrete implementation, threading model or other clients of any Web Bean it depends upon. Even better, the concrete implementation, lifecycle and threading model of a Web Bean it depends upon may vary according to the deployment scenario, without affecting any client.

Events, interceptors and decorators enhance the *loose-coupling* that is inherent in this model:

- *event notifications* decouple event producers from event consumers,
 - *interceptors* decouple technical concerns from business logic, and
-

- *decorators* allow business concerns to be compartmentalized.

Most importantly, Web Beans provides all these facilities in a *typesafe* way. Web Beans never uses string-based identifiers to determine how collaborating objects fit together. And XML, though it remains an option, is rarely used. Instead, Web Beans uses the typing information that is already available in the Java object model, together with a new pattern, called *binding annotations*, to wire together Web Beans, their dependencies, their interceptors and decorators and their event consumers.

The Web Beans services are general and apply to the following types of components that exist in the Java EE environment:

- all JavaBeans,
- all EJBs, and
- all Servlets.

Web Beans even provides the necessary integration points so that other kinds of components defined by future Java EE specifications or by non-standard frameworks may be cleanly integrated with Web Beans, take advantage of the Web Beans services, and interact with any other kind of Web Bean.

Web Beans was influenced by a number of existing Java frameworks, including Seam, Guice and Spring. However, Web Beans has its own very distinct character: more typesafe than Seam, more stateful and less XML-centric than Spring, more web and enterprise-application capable than Guice.

Most importantly, Web Beans is a JCP standard that integrates cleanly with Java EE, and with any Java SE environment where embeddable EJB Lite is available.

Getting started with Web Beans

So you're already keen to get started writing your first Web Bean? Or perhaps you're skeptical, wondering what kinds of hoops the Web Beans specification will make you jump through! The good news is that you've probably already written and used hundreds, perhaps thousands of Web Beans. You might not even remember the first Web Bean you wrote.

1.1. Your first Web Bean

With certain, very special exceptions, every Java class with a constructor that accepts no parameters is a Web Bean. That includes every `JavaBean`. Furthermore, every EJB 3-style session bean is a Web Bean. Sure, the `JavaBeans` and `EJBs` you've been writing every day have not been able to take advantage of the new services defined by the Web Beans specification, but you'll be able to use every one of them as Web Beans # injecting them into other Web Beans, configuring them via the Web Beans XML configuration facility, even adding interceptors and decorators to them # without touching your existing code.

Suppose that we have two existing Java classes, that we've been using for years in various applications. The first class parses a string into a list of sentences:

```
public class SentenceParser {  
    public List<String> parse(String text) { ... }  
}
```

The second existing class is a stateless session bean front-end for an external system that is able to translate sentences from one language to another:

```
@Stateless  
public class SentenceTranslator implements Translator {  
    public String translate(String sentence) { ... }  
}
```

Where `Translator` is the local interface:

```
@Local  
public interface Translator {  
    public String translate(String sentence);  
}
```

Unfortunately, we don't have a preexisting class that translates whole text documents. So let's write a Web Bean that does this job:

```
public class TextTranslator {

    private SentenceParser sentenceParser;
    private Translator sentenceTranslator;

    @Initializer
    TextTranslator(SentenceParser sentenceParser, Translator sentenceTranslator) {
        this.sentenceParser = sentenceParser;
        this.sentenceTranslator = sentenceTranslator;
    }

    public String translate(String text) {
        StringBuilder sb = new StringBuilder();
        for (String sentence: sentenceParser.parse(text)) {
            sb.append(sentenceTranslator.translate(sentence));
        }
        return sb.toString();
    }

}
```

We may obtain an instance of `TextTranslator` by injecting it into a Web Bean, Servlet or EJB:

```
@Initializer
public setTextTranslator(TextTranslator textTranslator) {
    this.textTranslator = textTranslator;
}
```

Alternatively, we may obtain an instance by directly calling a method of the Web Bean manager:

```
TextTranslator tt = manager.getInstanceByType(TextTranslator.class);
```

But wait: `TextTranslator` does not have a constructor with no parameters! Is it still a Web Bean? Well, a class that does not have a constructor with no parameters can still be a Web Bean if it has a constructor annotated `@Initializer`.

As you've guessed, the `@Initializer` annotation has something to do with dependency injection! `@Initializer` may be applied to a constructor or method of a Web Bean, and tells the Web Bean

manager to call that constructor or method when instantiating the Web Bean. The Web Bean manager will inject other Web Beans to the parameters of the constructor or method.

At system initialization time, the Web Bean manager must validate that exactly one Web Bean exists which satisfies each injection point. In our example, if no implementation of `Translator` available # if the `SentenceTranslator` EJB was not deployed # the Web Bean manager would throw an `UnsatisfiedDependencyException`. If more than one implementation of `Translator` was available, the Web Bean manager would throw an `AmbiguousDependencyException`.

1.2. What is a Web Bean?

So what, *exactly*, is a Web Bean?

A Web Bean is an application class that contains business logic. A Web Bean may be called directly from Java code, or it may be invoked via Unified EL. A Web Bean may access transactional resources. Dependencies between Web Beans are managed automatically by the Web Bean manager. Most Web Beans are *stateful* and *contextual*. The lifecycle of a Web Bean is always managed by the Web Bean manager.

Let's back up a second. What does it really mean to be "contextual"? Since Web Beans may be stateful, it matters *which* bean instance I have. Unlike a stateless component model (for example, stateless session beans) or a singleton component model (such as servlets, or singleton beans), different clients of a Web Bean see the Web Bean in different states. The client-visible state depends upon which instance of the Web Bean the client has a reference to.

However, like a stateless or singleton model, but *unlike* stateful session beans, the client does not control the lifecycle of the instance by explicitly creating and destroying it. Instead, the *scope* of the Web Bean determines:

- the lifecycle of each instance of the Web Bean and
- which clients share a reference to a particular instance of the Web Bean.

For a given thread in a Web Beans application, there may be an *active context* associated with the scope of the Web Bean. This context may be unique to the thread (for example, if the Web Bean is request scoped), or it may be shared with certain other threads (for example, if the Web Bean is session scoped) or even all other threads (if it is application scoped).

Clients (for example, other Web Beans) executing in the same context will see the same instance of the Web Bean. But clients in a different context will see a different instance.

One great advantage of the contextual model is that it allows stateful Web Beans to be treated like services! The client need not concern itself with managing the lifecycle of the Web Bean it is using, *nor does it even need to know what that lifecycle is*. Web Beans interact by passing messages, and the Web Bean implementations define the lifecycle of their own state. The Web Beans are loosely coupled because:

- they interact via well-defined public APIs

- their lifecycles are completely decoupled

We can replace one Web Bean with a different Web Bean that implements the same API and has a different lifecycle (a different scope) without affecting the other Web Bean implementation. In fact, Web Beans defines a sophisticated facility for overriding Web Bean implementations at deployment time, as we will see in [Section 4.2, “Deployment types”](#).

Note that not all clients of a Web Bean are Web Beans. Other objects such as Servlets or Message-Driven Beans # which are by nature not injectable, contextual objects # may also obtain references to Web Beans by injection.

Enough hand-waving. More formally, according to the spec:

A Web Bean comprises:

- A (nonempty) set of API types
- A (nonempty) set of binding annotation types
- A scope
- A deployment type
- Optionally, a Web Bean name
- A set of interceptor binding types
- A Web Bean implementation

Let's see what some of these terms mean, to the Web Bean developer.

1.2.1. API types, binding types and dependency injection

Web Beans usually acquire references to other Web Beans via dependency injection. Any injected attribute specifies a "contract" that must be satisfied by the Web Bean to be injected. The contract is:

- an API type, together with
- a set of binding types.

An API is a user-defined class or interface. (If the Web Bean is an EJB session bean, the API type is the `@Local` interface or bean-class local view). A binding type represents some client-visible semantic that is satisfied by some implementations of the API and not by others.

Binding types are represented by user-defined annotations that are themselves annotated `@BindingType`. For example, the following injection point has API type `PaymentProcessor` and binding type `@CreditCard`:

```
@CreditCard PaymentProcessor paymentProcessor
```

If no binding type is explicitly specified at an injection point, the default binding type `@Current` is assumed.

For each injection point, the Web Bean manager searches for a Web Bean which satisfies the contract (implements the API, and has all the binding types), and injects that Web Bean.

The following Web Bean has the binding type `@CreditCard` and implements the API type `PaymentProcessor`. It could therefore be injected to the example injection point:

```
@CreditCard
public class CreditCardPaymentProcessor
    implements PaymentProcessor { ... }
```

If a Web Bean does not explicitly specify a set of binding types, it has exactly one binding type: the default binding type `@Current`.

Web Beans defines a sophisticated but intuitive *resolution algorithm* that helps the container decide what to do if there is more than one Web Bean that satisfies a particular contract. We'll get into the details in [Chapter 4, Dependency injection](#).

1.2.2. Deployment types

Deployment types let us classify our Web Beans by deployment scenario. A deployment type is an annotation that represents a particular deployment scenario, for example `@Mock`, `@Staging` or `@AustralianTaxLaw`. We apply the annotation to Web Beans which should be deployed in that scenario. A deployment type allows a whole set of Web Beans to be conditionally deployed, with a just single line of configuration.

Many Web Beans just use the default deployment type `@Production`, in which case no deployment type need be explicitly specified. All three Web Bean in our example have the deployment type `@Production`.

In a testing environment, we might want to replace the `SentenceTranslator` Web Bean with a "mock object":

```
@Mock
public class MockSentenceTranslator implements Translator {
    public String translate(String sentence) {
        return "Lorem ipsum dolor sit amet";
    }
}
```

We would enable the deployment type `@Mock` in our testing environment, to indicate that `MockSentenceTranslator` and any other Web Bean annotated `@Mock` should be used.

We'll talk more about this unique and powerful feature in [Section 4.2, "Deployment types"](#).

1.2.3. Scope

The *scope* defines the lifecycle and visibility of instances of the Web Bean. The Web Beans context model is extensible, accommodating arbitrary scopes. However, certain important scopes are built-in to the specification, and provided by the Web Bean manager. A scope is represented by an annotation type.

For example, any web application may have *session scoped* Web Beans:

```
@SessionScoped
public class ShoppingCart { ... }
```

An instance of a session scoped Web Bean is bound to a user session and is shared by all requests that execute in the context of that session.

By default, Web Beans belong to a special scope called the *dependent pseudo-scope*. Web Beans with this scope are pure dependent objects of the object into which they are injected, and their lifecycle is bound to the lifecycle of that object.

We'll talk more about scopes in [Chapter 5, Scopes and contexts](#).

1.2.4. Web Bean names and Unified EL

A Web Bean may have a *name*, allowing it to be used in Unified EL expressions. It's easy to specify the name of a Web Bean:

```
@SessionScoped @Named("cart")
public class ShoppingCart { ... }
```

Now we can easily use the Web Bean in any JSF or JSP page:

```
<h:dataTable value="#{cart.lineItems}" var="item">
  ...
</h:dataTable>
```

It's even easier to just let the name be defaulted by the Web Bean manager:

```
@SessionScoped @Named
public class ShoppingCart { ... }
```


In this case, the name defaults to `shoppingCart` # the unqualified class name, with the first character changed to lowercase.

1.2.5. Interceptor binding types

Web Beans supports the interceptor functionality defined by EJB 3, not only for EJB beans, but also for plain Java classes. In addition, Web Beans provides a new approach to binding interceptors to EJB beans and other Web Beans.

It remains possible to directly specify the interceptor class via use of the `@Interceptors` annotation:

```
@SessionScoped
@Interceptors(TransactionInterceptor.class)
public class ShoppingCart { ... }
```

However, it is more elegant, and better practice, to indirect the interceptor binding through an *interceptor binding type*:

```
@SessionScoped @Transactional
public class ShoppingCart { ... }
```

We'll discuss Web Beans interceptors and decorators in [Chapter 7, Interceptors](#) and [Chapter 8, Decorators](#).

1.3. What kinds of objects can be Web Beans?

We've already seen that JavaBeans, EJBs and some other Java classes can be Web Beans. But exactly what kinds of objects are Web Beans?

1.3.1. Simple Web Beans

The Web Beans specification says that a concrete Java class is a *simple* Web Bean if:

- it is not an EE container-managed component, like an EJB, a Servlet or a JPA entity,
- it is not a non-static static inner class,
- it is not a parameterized type, and
- it has a constructor with no parameters, or a constructor annotated `@Initializer`.

Thus, almost every JavaBean is a simple Web Bean.

Every interface implemented directly or indirectly by a simple Web Bean is an API type of the simple Web Bean. The class and its superclasses are also API types.

1.3.2. Enterprise Web Beans

The specification says that all EJB 3-style session and singleton beans are *enterprise* Web Beans. Message driven beans are not Web Beans # since they are not intended to be injected into other objects # but they can take advantage of most of the functionality of Web Beans, including dependency injection and interceptors.

Every local interface of an enterprise Web Bean that does not have a wildcard type parameter or type variable, and every one of its superinterfaces, is an API type of the enterprise Web Bean. If the EJB bean has a bean class local view, the bean class, and every one of its superclasses, is also an API type.

Stateful session beans should declare a remove method with no parameters or a remove method annotated `@Destructor`. The Web Bean manager calls this method to destroy the stateful session bean instance at the end of its lifecycle. This method is called the *destructor* method of the enterprise Web Bean.

```
@Stateful @SessionScoped
public class ShoppingCart {

    ...

    @Remove
    public void destroy() {}

}
```

So when should we use an enterprise Web Bean instead of a simple Web Bean? Well, whenever we need the advanced enterprise services offered by EJB, such as:

- method-level transaction management and security,
- concurrency management,
- instance-level passivation for stateful session beans and instance-pooling for stateless session beans,
- remote and web service invocation, and
- timers and asynchronous methods,

we should use an enterprise Web Bean. When we don't need any of these things, a simple Web Bean will serve just fine.

Many Web Beans (including any session or application scoped Web Bean) are available for concurrent access. Therefore, the concurrency management provided by EJB 3.1 is especially useful. Most session and application scoped Web Beans should be EJBs.

Web Beans which hold references to heavy-weight resources, or hold a lot of internal state benefit from the advanced container-managed lifecycle defined by the EJB `@Stateless/@Stateful/@Singleton` model, with its support for passivation and instance pooling.

Finally, it's usually obvious when method-level transaction management, method-level security, timers, remote methods or asynchronous methods are needed.

It's usually easy to start with simple Web Bean, and then turn it into an EJB, just by adding an annotation: `@Stateless`, `@Stateful` or `@Singleton`.

1.3.3. Producer methods

A *producer method* is a method that is called by the Web Bean manager to obtain an instance of the Web Bean when no instance exists in the current context. A producer method lets the application take full control of the instantiation process, instead of leaving instantiation to the Web Bean manager. For example:

```
@ApplicationScoped
public class Generator {

    private Random random = new Random( System.currentTimeMillis() );

    @Produces @Random int next() {
        return random.nextInt(100);
    }

}
```

The result of a producer method is injected just like any other Web Bean.

```
@Random int randomNumber
```

The method return type and all interfaces it extends/implements directly or indirectly are API types of the producer method. If the return type is a class, all superclasses are also API types.

Some producer methods return objects that require explicit destruction:

```
@Produces @RequestScoped Connection connect(User user) {
    return createConnection( user.getId(), user.getPassword() );
}
```

```
}
```

These producer methods may define matching *disposal methods*:

```
void close(@Disposes Connection connection) {  
    connection.close();  
}
```

This disposal method is called automatically by the Web Bean manager at the end of the request.

We'll talk much more about producer methods in [Chapter 6, *Producer methods*](#).

1.3.4. JMS endpoints

Finally, a JMS queue or topic can be a Web Bean. Web Beans relieves the developer from the tedium of managing the lifecycles of all the various JMS objects required to send messages to queues and topics. We'll discuss JMS endpoints in [Section 13.4, “JMS endpoints”](#).

JSF web application example

Let's illustrate these ideas with a full example. We're going to implement user login/logout for an application that uses JSF. First, we'll define a Web Bean to hold the username and password entered during login:

```
@Named @RequestScoped
public class Credentials {

    private String username;
    private String password;

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

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

}
```

This Web Bean is bound to the login prompt in the following JSF form:

```
<h:form>
  <h:panelGrid columns="2" rendered="#{!login.loggedIn}">
    <h:outputLabel for="username">Username:</h:outputLabel>
    <h:inputText id="username" value="#{credentials.username}"/>
    <h:outputLabel for="password">Password:</h:outputLabel>
    <h:inputText id="password" value="#{credentials.password}"/>
  </h:panelGrid>
  <h:commandButton value="Login" action="#{login.login}" rendered="#{!login.loggedIn}"/>
  <h:commandButton value="Logout" action="#{login.logout}" rendered="#{login.loggedIn}"/>
</h:form>
```

The actual work is done by a session scoped Web Bean that maintains information about the currently logged-in user and exposes the `User` entity to other Web Beans:

```
@SessionScoped @Named
public class Login {

    @Current Credentials credentials;
```

```
@PersistenceContext EntityManager userDatabase;

private User user;

public void login() {

    List<User> results = userDatabase.createQuery(
        "select u from User u where u.username=:username and u.password=:password")
        .setParameter("username", credentials.getUsername())
        .setParameter("password", credentials.getPassword())
        .getResultList();

    if ( !results.isEmpty() ) {
        user = results.get(0);
    }

}

public void logout() {
    user = null;
}

public boolean isLoggedIn() {
    return user!=null;
}

@Produces @LoggedIn User getCurrentUser() {
    return user;
}

}
```

Of course, @LoggedIn is a binding annotation:

```
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD})
@BindingType
public @interface LoggedIn {}
```

Now, any other Web Bean can easily inject the current user:

```
public class DocumentEditor {
```

```
@Current Document document;  
@LoggedIn User currentUser;  
@PersistenceContext EntityManager docDatabase;  
  
public void save() {  
    document.setCreatedBy(currentUser);  
    docDatabase.persist(document);  
}  
  
}
```

Hopefully, this example gives a flavor of the Web Bean programming model. In the next chapter, we'll explore Web Beans dependency injection in greater depth.

The Web Beans Reference Implementation

The Web Beans Reference Implementation is being developed at [the Seam project](http://seamframework.org/WebBeans) [http://seamframework.org/WebBeans]. You can download the latest developer release of Web Beans from the [the downloads page](http://seamframework.org/Download) [http://seamframework.org/Download].

The Web Beans RI comes with a two deployable example applications: `webbeans-numberguess`, a war example, containing only simple beans, and `webbeans-translator` an ear example, containing enterprise beans. To run the examples you'll need the following:

- the latest release of the Web Beans RI,
- JBoss AS 5.0.0.GA, and
- Ant 1.7.0.

Currently, the Web Beans RI only runs on JBoss Application Server 5. You'll need to download JBoss AS 5.0.0.GA from [jboss.org](http://www.jboss.org/jbossas/downloads/) [http://www.jboss.org/jbossas/downloads/], and unzip it. For example:

```
$ cd /Applications
$ unzip ~/jboss-5.0.0.GA.zip
```

Next, download the Web Beans RI from [seamframework.org](http://seamframework.org/Download) [http://seamframework.org/Download], and unzip it. For example

```
$ cd ~/
$ unzip ~/webbeans-$VERSION.zip
```

Next, we need to tell Web Beans where JBoss is located. Edit `jboss-as/build.properties` and set the `jboss.home` property. For example:

```
jboss.home=/Applications/jboss-5.0.0.GA
```

As Web Beans is a new piece of software, you need to update JBoss AS to run the Web Beans RI. Future versions of JBoss AS will include these updates, and this step won't be necessary.



Note

Currently, two updates are needed. Firstly, a new deployer, `webbeans.deployer` is added. This adds supports for Web Bean archives to JBoss AS, and allows the Web Beans RI to query the EJB3 container and discover which EJBs are installed in your application. Secondly, an update to JBoss EJB3 is needed.

To install the update, you'll need Ant 1.7.0 installed, and the `ANT_HOME` environment variable set. For example:

```
$ unzip apache-ant-1.7.0.zip
$ export ANT_HOME=~/.apache-ant-1.7.0
```

Then, you can install the update. The update script will use Maven to download the Web Beans and EJB3 automatically.

```
$ cd webbeans-$VERSION/jboss-as
$ ant update
```

Now, you're ready to deploy your first example!



Tip

The build scripts for the examples offer a number of targets, these are:

- `ant restart` - deploy the example in exploded format
- `ant explode` - update an exploded example, without restarting the deployment
- `ant deploy` - deploy the example in compressed jar format
- `ant undeploy` - remove the example from the server
- `ant clean` - clean the example

To deploy the numberguess example:

```
$ cd examples/numberguess
ant deploy
```

Start JBoss AS:

```
$ /Application/jboss-5.0.0.GA/bin/run.sh
```



Tip

If you use Windows, use the `run.bat` script.

Wait for the application to deploy, and enjoy hours of fun at <http://localhost:8080/webbeans-numberguess!>

The Web Beans RI includes a second simple example that will translate your text into Latin. The numberguess example is a war example, and uses only simple beans; the translator example is an ear example, and includes enterprise beans, packaged in an EJB module. To try it out:

```
$ cd examples/translator
ant deploy
```

Wait for the application to deploy, and visit <http://localhost:8080/webbeans-translator!>

3.1. The numberguess example

In the numberguess application you get given 10 attempts to guess a number between 1 and 100. After each attempt, you will be told whether you are too high, or too low.

The numberguess example is comprised of a number of Web Beans, configuration files, and Facelet JSF pages, packaged as a war. Let's start with the configuration files.

All the configuration files for this example are located in `WEB-INF/`, which is stored in `WebContent` in the source tree. First, we have `faces-config.xml`, in which we tell JSF to use Facelets:

```
<?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/web-facesconfig_1_2.xsd">

  <application>
    <view-handler>com.sun.facelets.FaceletViewHandler</view-handler>
  </application>
```

```
</faces-config>
```

There is an empty `web-beans.xml` file, which marks this application as a Web Beans application.

Finally there is `web.xml`:

```
<?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">

  <display-name>Web Beans Numbergues example</display-name>

  <!-- JSF -->

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

  <servlet-mapping>                        2
    <servlet-name>Faces Servlet</servlet-name>
    <url-pattern>*.jsf</url-pattern>
  </servlet-mapping>

  <context-param>                          3
    <param-name>javax.faces.DEFAULT_SUFFIX</param-name>
    <param-value>.xhtml</param-value>
  </context-param>

  <session-config>                         4
    <session-timeout>10</session-timeout>
  </session-config>

</web-app>
```

- ① Enable and load the JSF servlet
- ② Configure requests to `.jsf` pages to be handled by JSF
- ③ Tell JSF that we will be giving our source files (facelets) an extension of `.jsf`
- ④ Configure a session timeout of 10 minutes



Note

Whilst this demo is a JSF demo, you can use the Web Beans RI with any Servlet based web framework.

Let's take a look at the Facelet view:

```
<!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: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">

  <ui:composition template="template.xhtml">                                ①
    <ui:define name="content">
      <h1>Guess a number...</h1>
      <h:form id="NumberGuessMain">

        <div style="color: red">                                           ②
          <h:messages id="messages" globalOnly="false"/>
          <h:outputText id="Higher" value="Higher!" rendered="#{game.number gt game.guess
and game.guess ne 0}"/>
          <h:outputText id="Lower" value="Lower!" rendered="#{game.number lt game.guess and
game.guess ne 0}"/>
        </div>

        <div>
          I'm thinking of a number between #{game.smallest} and #{game.bigge           ③
st}.
          You have #{game.remainingGuesses} guesses.
        </div>

        <div>
          Your guess:

          <h:inputText id="inputGuess"                                     ④

```

```
        value="#{game.guess}"
        required="true"
        size="3"
        disabled="#{game.number eq game.guess}">

        <f:validateLongRange maximum="#{game.biggest}"
            minimum="#{game.smallest}"/>
    </h:inputText>

    <h:commandButton id="GuessButton"
        value="Guess"
        action="#{game.check}"
        disabled="#{game.number eq game.guess}"/>
</div>
<div>
    <h:commandButton id="RestartButton" value="Reset" action="#{game.reset}"
immediate="true" />
</div>
</h:form>
</ui:define>
</ui:composition>
</html>
```

- ❶ Facelets is a templating language for JSF, here we are wrapping our page in a template which defines the header.
- ❷ There are a number of messages which can be sent to the user, "Higher!", "Lower!" and "Correct!"
- ❸ As the user guesses, the range of numbers they can guess gets smaller - this sentence changes to make sure they know what range to guess in.
- ❹ This input field is bound to a Web Bean, using the value expression.
- ❺ A range validator is used to make sure the user doesn't accidentally input a number outside of the range in which they can guess - if the validator wasn't here, the user might use up a guess on an out of range number.
- ❻ And, of course, there must be a way for the user to send their guess to the server. Here we bind to an action method on the Web Bean.

The example exists of 4 classes, the first two of which are binding types. First, there is the `@Random` binding type, used for injecting a random number:

```
@Target( { TYPE, METHOD, PARAMETER, FIELD })
@Retention(RUNTIME)
@Documented
@BindingType
```

```
public @interface Random {}
```

There is also the `@MaxNumber` binding type, used for injecting the maximum number that can be injected:

```
@Target( { TYPE, METHOD, PARAMETER, FIELD })
@Retention(RUNTIME)
@Documented
@BindingType
public @interface MaxNumber {}
```

The `Generator` class is responsible for creating the random number, via a producer method. It also exposes the maximum possible number via a producer method:

```
@ApplicationScoped
public class Generator {

    private java.util.Random random = new java.util.Random( System.currentTimeMillis() );

    private int maxNumber = 100;

    java.util.Random getRandom()
    {
        return random;
    }

    @Produces @Random int next() {
        return getRandom().nextInt(maxNumber);
    }

    @Produces @MaxNumber int getMaxNumber()
    {
        return maxNumber;
    }
}
```

You'll notice that the `Generator` is application scoped; therefore we don't get a different random each time.

The final Web Bean in the application is the session scoped `Game`.

You'll note that we've used the `@Named` annotation, so that we can use the bean through EL in the JSF page. Finally, we've used constructor injection to initialize the game with a random number. And of course, we need to tell the player when they've won, so we give feedback with a `FacesMessage`.

```
package org.jboss.webbeans.examples.numberguess;
```

```
import javax.annotation.PostConstruct;
import javax.faces.application.FacesMessage;
import javax.faces.context.FacesContext;
import javax.webbeans.AnnotationLiteral;
import javax.webbeans.Current;
import javax.webbeans.Initializer;
import javax.webbeans.Named;
import javax.webbeans.SessionScoped;
import javax.webbeans.manager.Manager;
```

```
@Named
```

```
@SessionScoped
```

```
public class Game
```

```
{
```

```
    private int number;
```

```
    private int guess;
```

```
    private int smallest;
```

```
    private int biggest;
```

```
    private int remainingGuesses;
```

```
    @Current Manager manager;
```

```
    public Game()
```

```
    {
```

```
    }
```

```
    @Initializer
```

```
    Game(@MaxNumber int maxNumber)
```

```
    {
```

```
        this.biggest = maxNumber;
```

```
    }
```

```
    public int getNumber()
```

```
    {
```

```
        return number;
```



```
}

public int getGuess()
{
    return guess;
}

public void setGuess(int guess)
{
    this.guess = guess;
}

public int getSmallest()
{
    return smallest;
}

public int getBiggest()
{
    return biggest;
}

public int getRemainingGuesses()
{
    return remainingGuesses;
}

public String check()
{
    if (guess > number)
    {
        biggest = guess - 1;
    }
    if (guess < number)
    {
        smallest = guess + 1;
    }
    if (guess == number)
    {
        FacesContext.getCurrentInstance().addMessage(null, new FacesMessage("Correct!"));
    }
    remainingGuesses--;
    return null;
}
```

```
@PostConstruct
public void reset()
{
    this.smallest = 0;
    this.guess = 0;
    this.remainingGuesses = 10;
    this.number = manager.getInstanceByType(Integer.class, new
AnnotationLiteral<Random>({});
}
}
```

3.2. The translator example

The translator example will take any sentences you enter, and translate them to Latin.

The translator example is built as an ear, and contains EJBs. As a result, it's structure is more complex than the numberguess example.



Note

EJB 3.1 and Jave EE 6 allow you to package EJBs in a war, which will make this structure much simpler!

First, let's take a look at the ear aggregator, which is located in `webbeans-translator-ear` module. Maven automatically generates the `application.xml` for us:

```
<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-ear-plugin</artifactId>
  <configuration>
    <modules>
      <webModule>
        <groupId>org.jboss.webbeans.examples.translator</groupId>
        <artifactId>webbeans-translator-war</artifactId>
        <contextRoot>/webbeans-translator</contextRoot>
      </webModule>
    </modules>
  </configuration>
</plugin>
```

Here we set the context path, which gives us a nice url (<http://localhost:8080/webbeans-translator>).



Tip

If you aren't using Maven to generate these files, you would need `META-INF/application.xml`:

```
<?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>webbeans-translator-ear</display-name>
    <description>Ear Example for the reference implementation of JSR 299: Web
Beans</description>

    <module>
        <web>
            <web-uri>webbeans-translator.war</web-uri>
            <context-root>/webbeans-translator</context-root>
        </web>
    </module>
    <module>
        <ejb>webbeans-translator.jar</ejb>
    </module>
</application>
```

Next, let's look at the war. Just as in the numberguess example, we have a `faces-config.xml` (to enable Facelets) and a `web.xml` (to enable JSF) in `WebContent/WEB-INF`.

More interesting is the facelet used to translate text. Just as in the numberguess example we have a template, which surrounds the form (omitted here for brevity):

```
<h:form id="NumberGuessMain">

    <table>
        <tr align="center" style="font-weight: bold" >
            <td>
                Your text
            </td>
```

```
<td>
    Translation
</td>
</tr>
<tr>
<td>
<h:inputTextarea id="text" value="#{translator.text}" required="true" rows="5" cols="80" />
</td>
<td>
<h:outputText value="#{translator.translatedText}" />
</td>
</tr>
</table>
<div>
<h:commandButton id="button" value="Translate" action="#{translator.translate}" />
</div>

</h:form>
```

The user can enter some text in the lefthand textarea, and hit the translate button to see the result to the right.

Finally, let's look at the ejb module, `webbeans-translator-ejb`. In `src/main/resources/META-INF` there is just an empty `web-beans.xml`, used to mark the archive as containing Web Beans.

We've saved the most interesting bit to last, the code! The project has two simple beans, `SentenceParser` and `TextTranslator` and two enterprise beans, `TranslatorControllerBean` and `SentenceTranslator`. You should be getting quite familiar with what a Web Bean looks like by now, so we'll just highlight the most interesting bits here.

Both `SentenceParser` and `TextTranslator` are dependent beans, and `TextTranslator` uses constructor initialization:

```
public class TextTranslator {
    private SentenceParser sentenceParser;
    private Translator sentenceTranslator;

    @Initializer
    TextTranslator(SentenceParser sentenceParser, Translator sentenceTranslator)
    {
        this.sentenceParser = sentenceParser;
        this.sentenceTranslator = sentenceTranslator;
    }
}
```

`TextTranslator` is a stateless bean (with a local business interface), where the magic happens - of course, we couldn't develop a full translator, but we gave it a good go!

Finally, there is UI orientated controller, that collects the text from the user, and dispatches it to the translator. This is a request scoped, named, stateful session bean, which injects the translator.

```
@Stateful
@RequestScoped
@Named("translator")
public class TranslatorControllerBean implements TranslatorController
{

    @Current TextTranslator translator;
```

The bean also has getters and setters for all the fields on the page.

As this is a stateful session bean, we have to have a remove method:

```
@Remove
public void remove()
{

}
```

The Web Beans manager will call the remove method for you when the bean is destroyed; in this case at the end of the request.

That concludes our short tour of the Web Beans RI examples. For more on the RI, or to help out, please visit <http://www.seamframework.org/WebBeans/Development>.

We need help in all areas - bug fixing, writing new features, writing examples and translating this reference guide.

Dependency injection

Web Beans supports three primary mechanisms for dependency injection:

Constructor parameter injection:

```
public class Checkout {  
  
    private final ShoppingCart cart;  
  
    @Initializer  
    public Checkout(ShoppingCart cart) {  
        this.cart = cart;  
    }  
  
}
```

Initializer method parameter injection:

```
public class Checkout {  
  
    private ShoppingCart cart;  
  
    @Initializer  
    void setShoppingCart(ShoppingCart cart) {  
        this.cart = cart;  
    }  
  
}
```

And direct field injection:

```
public class Checkout {  
  
    private @Current ShoppingCart cart;  
  
}
```

Dependency injection always occurs when the Web Bean instance is first instantiated.

- First, the Web Bean manager calls the Web Bean constructor, to obtain an instance of the Web Bean.
- Next, the Web Bean manager initializes the values of all injected fields of the Web Bean.
- Next, the Web Bean manager calls all initializer methods of Web Bean.
- Finally, the `@PostConstruct` method of the Web Bean, if any, is called.

Constructor parameter injection is not supported for EJB beans, since the EJB is instantiated by the EJB container, not the Web Bean manager.

Parameters of constructors and initializer methods need not be explicitly annotated when the default binding type `@Current` applies. Injected fields, however, *must* specify a binding type, even when the default binding type applies. If the field does not specify a binding type, it will not be injected.

Producer methods also support parameter injection:

```
@Produces Checkout createCheckout(ShoppingCart cart) {  
    return new Checkout(cart);  
}
```

Finally, observer methods (which we'll meet in [Chapter 9, Events](#)), disposal methods and destructor methods all support parameter injection.

The Web Beans specification defines a procedure, called the *typesafe resolution algorithm*, that the Web Bean manager follows when identifying the Web Bean to inject to an injection point. This algorithm looks complex at first, but once you understand it, it's really quite intuitive. Typesafe resolution is performed at system initialization time, which means that the manager will inform the user immediately if a Web Bean's dependencies cannot be satisfied, by throwing a `UnsatisfiedDependencyException` or `AmbiguousDependencyException`.

The purpose of this algorithm is to allow multiple Web Beans to implement the same API type and either:

- allow the client to select which implementation it requires using *binding annotations*,
- allow the application deployer to select which implementation is appropriate for a particular deployment, without changes to the client, by enabling or disabling *deployment types*, or
- allow one implementation of an API to override another implementation of the same API at deployment time, without changes to the client, using *deployment type precedence*.

Let's explore how the Web Beans manager determines a Web Bean to be injected.

4.1. Binding annotations

If we have more than one Web Bean that implements a particular API type, the injection point can specify exactly which Web Bean should be injected using a binding annotation. For example, there might be two implementations of `PaymentProcessor`:

```
@PayByCheque
public class ChequePaymentProcessor implements PaymentProcessor {
    public void process(Payment payment) { ... }
}
```

```
@PayByCreditCard
public class CreditCardPaymentProcessor implements PaymentProcessor {
    public void process(Payment payment) { ... }
}
```

Where `@PayByCheque` and `@PayByCreditCard` are binding annotations:

```
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD, PARAMETER})
@BindingType
public @interface PayByCheque {}
```

```
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD, PARAMETER})
@BindingType
public @interface PayByCreditCard {}
```

A client Web Bean developer uses the binding annotation to specify exactly which Web Bean should be injected.

Using field injection:

```
@PayByCheque PaymentProcessor chequePaymentProcessor;
@PayByCreditCard PaymentProcessor creditCardPaymentProcessor;
```

Using initializer method injection:

```
@Initializer
public void setPaymentProcessors(@PayByCheque PaymentProcessor
    chequePaymentProcessor,
    @PayByCreditCard PaymentProcessor creditCardPaymentProcessor) {
    this.chequePaymentProcessor = chequePaymentProcessor;
    this.creditCardPaymentProcessor = creditCardPaymentProcessor;
}
```

Or using constructor injection:

```
@Initializer
public Checkout(@PayByCheque PaymentProcessor chequePaymentProcessor,
    @PayByCreditCard PaymentProcessor creditCardPaymentProcessor) {
    this.chequePaymentProcessor = chequePaymentProcessor;
    this.creditCardPaymentProcessor = creditCardPaymentProcessor;
}
```

4.1.1. Binding annotations with members

Binding annotations may have members:

```
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD, PARAMETER})
@BindingType
public @interface PayBy {
    PaymentType value();
}
```

In which case, the member value is significant:

```
@PayBy(CHEQUE) PaymentProcessor chequePaymentProcessor;
@PayBy(CREDIT_CARD) PaymentProcessor creditCardPaymentProcessor;
```

You can tell the Web Bean manager to ignore a member of a binding annotation type by annotating the member `@NonBinding`.

4.1.2. Combinations of binding annotations

An injection point may even specify multiple binding annotations:

```
@Asynchronous @PayByCheque PaymentProcessor paymentProcessor
```

In this case, only a Web Bean which has *both* binding annotations would be eligible for injection.

4.1.3. Binding annotations and producer methods

Even producer methods may specify binding annotations:

```
@Produces
@Asynchronous @PayByCheque
PaymentProcessor createAsyncPaymentProcessor(@PayByCheque PaymentProcessor
processor) {
    return new AsynchronousPaymentProcessor(processor);
}
```

4.1.4. The default binding type

Web Beans defines a binding type `@Current` that is the default binding type for any injection point or Web Bean that does not explicitly specify a binding type.

There are two common circumstances in which it is necessary to explicitly specify `@Current`:

- on a field, in order to declare it as an injected field with the default binding type, and
- on a Web Bean which has another binding type in addition to the default binding type.

4.2. Deployment types

All Web Beans have a *deployment type*. Each deployment type identifies a set of Web Beans that should be conditionally installed in some deployments of the system.

For example, we could define a deployment type named `@Mock`, which would identify Web Beans that should only be installed when the system executes inside an integration testing environment:

```
@Retention(RUNTIME)
@Target({TYPE, METHOD})
@DeploymentType
public @interface Mock {}
```

Suppose we had some Web Bean that interacted with an external system to process payments:

```
public class ExternalPaymentProcessor {  
  
    public void process(Payment p) {  
        ...  
    }  
  
}
```

Since this Web Bean does not explicitly specify a deployment type, it has the default deployment type `@Production`.

For integration or unit testing, the external system is slow or unavailable. So we would create a mock object:

```
@Mock  
public class MockPaymentProcessor implements PaymentProcessor {  
  
    @Override  
    public void process(Payment p) {  
        p.setSuccessful(true);  
    }  
  
}
```

But how does the Web Bean manager determine which implementation to use in a particular deployment?

4.2.1. Enabling deployment types

Web Beans defines two built-in deployment types: `@Production` and `@Standard`. By default, only Web Beans with the built-in deployment types are enabled when the system is deployed. We can identify additional deployment types to be enabled in a particular deployment by listing them in `web-beans.xml`.

Going back to our example, when we deploy our integration tests, we want all our `@Mock` objects to be installed:

```
<WebBeans>  
  <Deploy>  
    <Standard/>  
    <Production/>  
    <test:Mock/>  
  </Deploy>  
</WebBeans>
```

```
</Deploy>
</WebBeans>
```

Now the Web Bean manager will identify and install all Web Beans annotated `@Production`, `@Standard` or `@Mock` at deployment time.

The deployment type `@Standard` is used only for certain special Web Beans defined by the Web Beans specification. We can't use it for our own Web Beans, and we can't disable it.

The deployment type `@Production` is the default deployment type for Web Beans which don't explicitly declare a deployment type, and may be disabled.

4.2.2. Deployment type precedence

If you've been paying attention, you're probably wondering how the Web Bean manager decides which implementation `# ExternalPaymentProcessor` or `MockPaymentProcessor` `#` to choose. Consider what happens when the manager encounters this injection point:

```
@Current PaymentProcessor paymentProcessor
```

There are now two Web Beans which satisfy the `PaymentProcessor` contract. Of course, we can't use a binding annotation to disambiguate, since binding annotations are hard-coded into the source at the injection point, and we want the manager to be able to decide at deployment time!

The solution to this problem is that each deployment type has a different *precedence*. The precedence of the deployment types is determined by the order in which they appear in `web-beans.xml`. In our example, `@Mock` appears later than `@Production` so it has a higher precedence.

Whenever the manager discovers that more than one Web Bean could satisfy the contract (API type plus binding annotations) specified by an injection point, it considers the relative precedence of the Web Beans. If one has a higher precedence than the others, it chooses the higher precedence Web Bean to inject. So, in our example, the Web Bean manager will inject `MockPaymentProcessor` when executing in our integration testing environment (which is exactly what we want).

It's interesting to compare this facility to today's popular manager architectures. Various "lightweight" containers also allow conditional deployment of classes that exist in the classpath, but the classes that are to be deployed must be explicitly, individually, listed in configuration code or in some XML configuration file. Web Beans does support Web Bean definition and configuration via XML, but in the common case where no complex configuration is required, deployment types allow a whole set of Web Beans to be enabled with a single line of XML. Meanwhile, a developer browsing the code can easily identify what deployment scenarios the Web Bean will be used in.

4.2.3. Example deployment types

Deployment types are useful for all kinds of things, here's some examples:

- `@Mock` and `@Staging` deployment types for testing
- `@AustralianTaxLaw` for site-specific Web Beans
- `@SeamFramework`, `@Guice` for third-party frameworks which build on Web Beans
- `@Standard` for standard Web Beans defined by the Web Beans specification

I'm sure you can think of more applications...

4.3. Fixing unsatisfied dependencies

The typesafe resolution algorithm fails when, after considering the binding annotations and deployment types of all Web Beans that implement the API type of an injection point, the Web Bean manager is unable to identify exactly one Web Bean to inject.

It's usually easy to fix an `UnsatisfiedDependencyException` or `AmbiguousDependencyException`.

To fix an `UnsatisfiedDependencyException`, simply provide a Web Bean which implements the API type and has the binding types of the injection point # or enable the deployment type of a Web Bean that already implements the API type and has the binding types.

To fix an `AmbiguousDependencyException`, introduce a binding type to distinguish between the two implementations of the API type, or change the deployment type of one of the implementations so that the Web Bean manager can use deployment type precedence to choose between them. An `AmbiguousDependencyException` can only occur if two Web Beans share a binding type and have exactly the same deployment type.

There's one more issue you need to be aware of when using dependency injection in Web Beans.

4.4. Client proxies

Clients of an injected Web Bean do not usually hold a direct reference to a Web Bean instance.

Imagine that a Web Bean bound to the application scope held a direct reference to a Web Bean bound to the request scope. The application scoped Web Bean is shared between many different requests. However, each request should see a different instance of the request scoped Web bean!

Now imagine that a Web Bean bound to the session scope held a direct reference to a Web Bean bound to the application scope. From time to time, the session context is serialized to disk in order to use memory more efficiently. However, the application scoped Web Bean instance should not be serialized along with the session scoped Web Bean!

Therefore, unless a Web Bean has the default scope `@Dependent`, the Web Bean manager must indirect all injected references to the Web Bean through a proxy object. This *client proxy* is responsible for ensuring that the Web Bean instance that receives a method invocation is the instance that is associated with the current context. The client proxy also allows Web Beans bound to contexts such as the session context to be serialized to disk without recursively serializing other injected Web Beans.

Unfortunately, due to limitations of the Java language, some Java types cannot be proxied by the Web Bean manager. Therefore, the Web Bean manager throws an `UnproxyableDependencyException` if the type of an injection point cannot be proxied.

The following Java types cannot be proxied by the Web Bean manager:

- classes which are declared `final` or have a `final` method,
- classes which have no non-private constructor with no parameters, and
- arrays and primitive types.

It's usually very easy to fix an `UnproxyableDependencyException`. Simply add a constructor with no parameters to the injected class, introduce an interface, or change the scope of the injected Web Bean to `@Dependent`.

4.5. Obtaining a Web Bean by programatic lookup

The application may obtain an instance of the interface `Manager` by injection:

```
@Current Manager manager;
```

The `Manager` object provides a set of methods for obtaining a Web Bean instance programatically.

```
PaymentProcessor p = manager.getInstanceByType(PaymentProcessor.class);
```

Binding annotations may be specified by subclassing the helper class `AnnotationLiteral`, since it is otherwise difficult to instantiate an annotation type in Java.

```
PaymentProcessor p = manager.getInstanceByType(PaymentProcessor.class,  
new AnnotationLiteral<CreditCard>({});
```

If the binding type has an annotation member, we can't use an anonymous subclass of `AnnotationLiteral` # instead we'll need to create a named subclass:

```
abstract class CreditCardBinding  
    extends AnnotationLiteral<CreditCard>  
    implements CreditCard {}
```

```
PaymentProcessor p = manager.getInstanceByType(PaymentProcessor.class,  
new CreditCardBinding() {
```

```
public void value() { return paymentType; }  
});
```

4.6. Lifecycle callbacks, `@Resource`, `@EJB` and `@PersistenceContext`

Enterprise Web Beans support all the lifecycle callbacks defined by the EJB specification: `@PostConstruct`, `@PreDestroy`, `@PrePassivate` and `@PostActivate`.

Simple Web Beans support only the `@PostConstruct` and `@PreDestroy` callbacks.

Both enterprise and simple Web Beans support the use of `@Resource`, `@EJB` and `@PersistenceContext` for injection of Java EE resources, EJBs and JPA persistence contexts, respectively. Simple Web Beans do not support the use of `@PersistenceContext(type=EXTENDED)`.

The `@PostConstruct` callback always occurs after all dependencies have been injected.

4.7. The `InjectionPoint` object

There are certain kinds of dependent objects # Web Beans with scope `@Dependent` # that need to know something about the object or injection point into which they are injected in order to be able to do what they do. For example:

- The log category for a `Logger` depends upon the class of the object that owns it.
- Injection of a HTTP parameter or header value depends upon what parameter or header name was specified at the injection point.
- Injection of the result of an EL expression evaluation depends upon the expression that was specified at the injection point.

A Web Bean with scope `@Dependent` may inject an instance of `InjectionPoint` and access metadata relating to the injection point to which it belongs.

Let's look at an example. The following code is verbose, and vulnerable to refactoring problems:

```
Logger log = Logger.getLogger(MyClass.class.getName());
```

This clever little producer method lets you inject a JDK `Logger` without explicitly specifying the log category:

```
class LogFactory {  
  
    @Produces Logger createLogger(InjectionPoint injectionPoint) {  
        return Logger.getLogger(injectionPoint.getMember().getDeclaringClass().getName());  
    }  
}
```



```
}

}
```

We can now write:

```
@Current Logger log;
```

Not convinced? Then here's a second example. To inject HTTP parameters, we need to define a binding type:

```
@BindingType
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD, PARAMETER})
public @interface HttpParam {
    @NonBinding public String value();
}
```

We would use this binding type at injection points as follows:

```
@HttpParam("username") String username;
@HttpParam("password") String password;
```

The following producer method does the work:

```
class HttpParams

    @Produces @HttpParam("")
    String getParamValue(ServletRequest request, InjectionPoint ip) {
        return request.getParameter(ip.getAnnotation(HttpParam.class).value());
    }

}
```

(Note that the `value()` member of the `HttpParam` annotation is ignored by the Web Bean manager since it is annotated `@NonBinding`.)

The Web Bean manager provides a built-in Web Bean that implements the `InjectionPoint` interface:

```
public interface InjectionPoint {  
    public Object getInstance();  
    public Bean<?> getBean();  
    public Member getMember():  
    public <T extends Annotation> T getAnnotation(Class<T> annotation);  
    public Set<T extends Annotation> getAnnotations();  
}
```

Scopes and contexts

So far, we've seen a few examples of *scope type annotations*. The scope of a Web Bean determines the lifecycle of instances of the Web Bean. The scope also determines which clients refer to which instances of the Web Bean. According to the Web Beans specification, a scope determines:

- When a new instance of any Web Bean with that scope is created
- When an existing instance of any Web Bean with that scope is destroyed
- Which injected references refer to any instance of a Web Bean with that scope

For example, if we have a session scoped Web Bean, `CurrentUser`, all Web Beans that are called in the context of the same `HttpSession` will see the same instance of `CurrentUser`. This instance will be automatically created the first time a `CurrentUser` is needed in that session, and automatically destroyed when the session ends.

5.1. Scope types

Web Beans features an *extensible context model*. It is possible to define new scopes by creating a new scope type annotation:

```
@Retention(RUNTIME)
@Target({TYPE, METHOD})
@ScopeType
public @interface ClusterScoped {}
```

Of course, that's the easy part of the job. For this scope type to be useful, we will also need to define a `Context` object that implements the scope! Implementing a `Context` is usually a very technical task, intended for framework development only.

We can apply a scope type annotation to a Web Bean implementation class to specify the scope of the Web Bean:

```
@ClusterScoped
public class SecondLevelCache { ... }
```

Usually, you'll use one of Web Beans' built-in scopes.

5.2. Built-in scopes

Web Beans defines four built-in scopes:

- `@RequestScoped`
- `@SessionScoped`
- `@ApplicationScoped`
- `@ConversationScoped`

For a web application that uses Web Beans:

- any servlet request has access to active request, session and application scopes, and, additionally
- any JSF request has access to an active conversation scope.

The request and application scopes are also active:

- during invocations of EJB remote methods,
- during EJB timeouts,
- during message delivery to a message-driven bean, and
- during web service invocations.

If the application tries to invoke a Web Bean with a scope that does not have an active context, a `ContextNotActiveException` is thrown by the Web Bean manager at runtime.

Three of the four built-in scopes should be extremely familiar to every Java EE developer, so let's not waste time discussing them here. One of the scopes, however, is new.

5.3. The conversation scope

The Web Beans conversation scope is a bit like the traditional session scope in that it holds state associated with a user of the system, and spans multiple requests to the server. However, unlike the session scope, the conversation scope:

- is demarcated explicitly by the application, and
- holds state associated with a particular web browser tab in a JSF application.

A conversation represents a task, a unit of work from the point of view of the user. The conversation context holds state associated with what the user is currently working on. If the user is doing multiple things at the same time, there are multiple conversations.

The conversation context is active during any JSF request. However, most conversations are destroyed at the end of the request. If a conversation should hold state across multiple requests, it must be explicitly promoted to a *long-running conversation*.

5.3.1. Conversation demarcation

Web Beans provides a built-in Web Bean for controlling the lifecycle of conversations in a JSF application. This Web Bean may be obtained by injection:

```
@Current Conversation conversation;
```

To promote the conversation associated with the current request to a long-running conversation, call the `begin()` method from application code. To schedule the current long-running conversation context for destruction at the end of the current request, call `end()`.

In the following example, a conversation-scoped Web Bean controls the conversation with which it is associated:

```
@ConversationScoped @Stateful
public class OrderBuilder {

    private Order order;
    private @Current Conversation conversation;
    private @PersistenceContext(type=EXTENDED) EntityManager em;

    @Produces public Order getOrder() {
        return order;
    }

    public Order createOrder() {
        order = new Order();
        conversation.begin();
        return order;
    }

    public void addLineItem(Product product, int quantity) {
        order.add( new LineItem(product, quantity) );
    }

    public void saveOrder(Order order) {
        em.persist(order);
        conversation.end();
    }

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

```
}
```

This Web Bean is able to control its own lifecycle through use of the `Conversation` API. But some other Web Beans have a lifecycle which depends completely upon another object.

5.3.2. Conversation propagation

The conversation context automatically propagates with any JSF faces request (JSF form submission). It does not automatically propagate with non-faces requests, for example, navigation via a link.

We can force the conversation to propagate with a non-faces request by including the unique identifier of the conversation as a request parameter. The Web Beans specification reserves the request parameter named `cid` for this use. The unique identifier of the conversation may be obtained from the `Conversation` object, which has the Web Beans name `conversation`.

Therefore, the following link propagates the conversation:

```
<a href="/addProduct.jsp?cid=#{conversation.id}">Add Product</a>
```

The Web Bean manager is also required to propagate conversations across any redirect, even if the conversation is not marked long-running. This makes it very easy to implement the common POST-then-redirect pattern, without resort to fragile constructs such as a "flash" object. In this case, the Web Bean manager automatically adds a request parameter to the redirect URL.

5.3.3. Conversation timeout

The Web Bean manager is permitted to destroy a conversation and all state held in its context at any time in order to preserve resources. A Web Bean manager implementation will normally do this on the basis of some kind of timeout # though this is not required by the Web Beans specification. The timeout is the period of inactivity before the conversation is destroyed.

The `Conversation` object provides a method to set the timeout. This is a hint to the Web Bean manager, which is free to ignore the setting.

```
conversation.setTimeout(timeoutInMillis);
```

5.4. The dependent pseudo-scope

In addition to the four built-in scopes, Web Beans features the so-called *dependent pseudo-scope*. This is the default scope for a Web Bean which does not explicitly declare a scope type.

For example, this Web Bean has the scope type `@Dependent`:

```
public class Calculator { ... }
```

When an injection point of a Web Bean resolves to a dependent Web Bean, a new instance of the dependent Web Bean is created every time the first Web Bean is instantiated. Instances of dependent Web Beans are never shared between different Web Beans or different injection points. They are *dependent objects* of some other Web Bean instance.

Dependent Web Bean instances are destroyed when the instance they depend upon is destroyed.

Web Beans makes it easy to obtain a dependent instance of a Java class or EJB bean, even if the class or EJB bean is already declared as a Web Bean with some other scope type.

5.4.1. The @New annotation

The built-in @New binding annotation allows *implicit* definition of a dependent Web Bean at an injection point. Suppose we declare the following injected field:

```
@New Calculator calculator;
```

Then a Web Bean with scope @Dependent, binding type @New, API type Calculator, implementation class Calculator and deployment type @Standard is implicitly defined.

This is true even if Calculator is *already* declared with a different scope type, for example:

```
@ConversationScoped
public class Calculator { ... }
```

So the following injected attributes each get a different instance of Calculator:

```
public class PaymentCalc {

    @Current Calculator calculator;
    @New Calculator newCalculator;

}
```

The calculator field has a conversation-scoped instance of Calculator injected. The newCalculator field has a new instance of Calculator injected, with a lifecycle that is bound to the owning PaymentCalc.

This feature is particularly useful with producer methods, as we'll see in the next chapter.

Producer methods

Producer methods let us overcome certain limitations that arise when the Web Bean manager, instead of the application, is responsible for instantiating objects. They're also the easiest way to integrate objects which are not Web Beans into the Web Beans environment. (We'll meet a second approach in [Chapter 12, Defining Web Beans using XML](#).)

According to the spec:

A Web Beans producer method acts as a source of objects to be injected, where:

- the objects to be injected are not required to be instances of Web Beans,
- the concrete type of the objects to be injected may vary at runtime or
- the objects require some custom initialization that is not performed by the Web Bean constructor

For example, producer methods let us:

- expose a JPA entity as a Web Bean,
- expose any JDK class as a Web Bean,
- define multiple Web Beans, with different scopes or initialization, for the same implementation class, or
- vary the implementation of an API type at runtime.

In particular, producer methods let us use runtime polymorphism with Web Beans. As we've seen, deployment types are a powerful solution to the problem of deployment-time polymorphism. But once the system is deployed, the Web Bean implementation is fixed. A producer method has no such limitation:

```
@SessionScoped
public class Preferences {

    private PaymentStrategyType paymentStrategy;

    ...

    @Produces @Preferred
    public PaymentStrategy getPaymentStrategy() {
        switch (paymentStrategy) {
            case CREDIT_CARD: return new CreditCardPaymentStrategy();
            case CHEQUE: return new ChequePaymentStrategy();
            case PAYPAL: return new PayPalPaymentStrategy();
```

```
        default: return null;
    }
}
}
```

Consider an injection point:

```
@Preferred PaymentStrategy paymentStrat;
```

This injection point has the same type and binding annotations as the producer method, so it resolves to the producer method using the usual Web Beans injection rules. The producer method will be called by the Web Bean manager to obtain an instance to service this injection point.

6.1. Scope of a producer method

The scope of the producer method defaults to `@Dependent`, and so it will be called *every time* the Web Bean manager injects this field or any other field that resolves to the same producer method. Thus, there could be multiple instances of the `PaymentStrategy` object for each user session.

To change this behavior, we can add a `@SessionScoped` annotation to the method.

```
@Produces @Preferred @SessionScoped
public PaymentStrategy getPaymentStrategy() {
    ...
}
```

Now, when the producer method is called, the returned `PaymentStrategy` will be bound to the session context. The producer method won't be called again in the same session.

6.2. Injection into producer methods

There's one potential problem with the code above. The implementations of `CreditCardPaymentStrategy` are instantiated using the Java `new` operator. Objects instantiated directly by the application can't take advantage of dependency injection and don't have interceptors.

If this isn't what we want we can use dependency injection into the producer method to obtain Web Bean instances:

```
@Produces @Preferred @SessionScoped
public PaymentStrategy getPaymentStrategy(CreditCardPaymentStrategy ccps,
```

```

        ChequePaymentStrategy cps,
        PayPalPaymentStrategy ppps) {
    switch (paymentStrategy) {
        case CREDIT_CARD: return ccps;
        case CHEQUE: return cps;
        case PAYPAL: return ppps;
        default: return null;
    }
}

```

Wait, what if `CreditCardPaymentStrategy` is a request scoped Web Bean? Then the producer method has the effect of "promoting" the current request scoped instance into session scope. This is almost certainly a bug! The request scoped object will be destroyed by the Web Bean manager before the session ends, but the reference to the object will be left "hanging" in the session scope. This error will *not* be detected by the Web Bean manager, so please take extra care when returning Web Bean instances from producer methods!

There's at least three ways we could go about fixing this bug. We could change the scope of the `CreditCardPaymentStrategy` implementation, but this would affect other clients of that Web Bean. A better option would be to change the scope of the producer method to `@Dependent` or `@RequestScoped`.

But a more common solution is to use the special `@New` binding annotation.

6.3. Use of @New with producer methods

Consider the following producer method:

```

@Produces @Preferred @SessionScoped
public PaymentStrategy getPaymentStrategy(@New CreditCardPaymentStrategy ccps,
        @New ChequePaymentStrategy cps,
        @New PayPalPaymentStrategy ppps) {
    switch (paymentStrategy) {
        case CREDIT_CARD: return ccps;
        case CHEQUE: return cps;
        case PAYPAL: return ppps;
        default: return null;
    }
}

```

Then a new *dependent* instance of `CreditCardPaymentStrategy` will be created, passed to the producer method, returned by the producer method and finally bound to the session context. The dependent object won't be destroyed until the `Preferences` object is destroyed, at the end of the session.

Part II. Developing loosely-coupled code

The first major theme of Web Beans is *loose coupling*. We've already seen three means of achieving loose coupling:

- *deployment types* enable deployment time polymorphism,
- *producer methods* enable runtime polymorphism, and
- *contextual lifecycle management* decouples Web Bean lifecycles.

These techniques serve to enable loose coupling of client and server. The client is no longer tightly bound to an implementation of an API, nor is it required to manage the lifecycle of the server object. This approach lets *stateful objects interact as if they were services*.

Loose coupling makes a system more *dynamic*. The system can respond to change in a well-defined manner. In the past, frameworks that attempted to provide the facilities listed above invariably did it by sacrificing type safety. Web Beans is the first technology that achieves this level of loose coupling in a typesafe way.

Web Beans provides three extra important facilities that further the goal of loose coupling:

- *interceptors* decouple technical concerns from business logic,
- *decorators* may be used to decouple some business concerns, and
- *event notifications* decouple event producers from event consumers.

Let's explore interceptors first.

Interceptors

Web Beans re-uses the basic interceptor architecture of EJB 3.0, extending the functionality in two directions:

- Any Web Bean may have interceptors, not just session beans.
- Web Beans features a more sophisticated annotation-based approach to binding interceptors to Web Beans.

The EJB specification defines two kinds of interception points:

- business method interception, and
- lifecycle callback interception.

A *business method interceptor* applies to invocations of methods of the Web Bean by clients of the Web Bean:

```
public class TransactionInterceptor {  
    @AroundInvoke public Object manageTransaction(InvocationContext ctx) { ... }  
}
```

A *lifecycle callback interceptor* applies to invocations of lifecycle callbacks by the container:

```
public class DependencyInjectionInterceptor {  
    @PostConstruct public void injectDependencies(InvocationContext ctx) { ... }  
}
```

An interceptor class may intercept both lifecycle callbacks and business methods.

7.1. Interceptor bindings

Suppose we want to declare that some of our Web Beans are transactional. The first thing we need is an *interceptor binding annotation* to specify exactly which Web Beans we're interested in:

```
@InterceptorBindingType  
@Target({METHOD, TYPE})  
@Retention(RUNTIME)  
public @interface Transactional {}
```

Now we can easily specify that our `ShoppingCart` is a transactional object:

```
@Transactional
public class ShoppingCart { ... }
```

Or, if we prefer, we can specify that just one method is transactional:

```
public class ShoppingCart {
    @Transactional public void checkout() { ... }
}
```

7.2. Implementing interceptors

That's great, but somewhere along the line we're going to have to actually implement the interceptor that provides this transaction management aspect. All we need to do is create a standard EJB interceptor, and annotate it `@Interceptor` and `@Transactional`.

```
@Transactional @Interceptor
public class TransactionInterceptor {
    @AroundInvoke public Object manageTransaction(InvocationContext ctx) { ... }
}
```

All Web Beans interceptors are simple Web Beans, and can take advantage of dependency injection and contextual lifecycle management.

```
@ApplicationScoped @Transactional @Interceptor
public class TransactionInterceptor {

    @Resource Transaction transaction;

    @AroundInvoke public Object manageTransaction(InvocationContext ctx) { ... }

}
```

Multiple interceptors may use the same interceptor binding type.

7.3. Enabling interceptors

Finally, we need to *enable* our interceptor in `web-beans.xml`.


```
<Interceptors>
  <tx:TransactionInterceptor/>
</Interceptors>
```

Whoah! Why the angle bracket stew?

Well, the XML declaration solves two problems:

- it enables us to specify a total ordering for all the interceptors in our system, ensuring deterministic behavior, and
- it lets us enable or disable interceptor classes at deployment time.

For example, we could specify that our security interceptor runs before our `TransactionInterceptor`.

```
<Interceptors>
  <sx:SecurityInterceptor/>
  <tx:TransactionInterceptor/>
</Interceptors>
```

Or we could turn them both off in our test environment!

7.4. Interceptor bindings with members

Suppose we want to add some extra information to our `@Transactional` annotation:

```
@InterceptorBindingType
@Target({METHOD, TYPE})
@Retention(RUNTIME)
public @interface Transactional {
    boolean requiresNew() default false;
}
```

Web Beans will use the value of `requiresNew` to choose between two different interceptors, `TransactionInterceptor` and `RequiresNewTransactionInterceptor`.

```
@Transactional(requiresNew=true) @Interceptor
public class RequiresNewTransactionInterceptor {
    @AroundInvoke public Object manageTransaction(InvocationContext ctx) { ... }
```

```
}
```

Now we can use `RequiresNewTransactionInterceptor` like this:

```
@Transactional(requiresNew=true)
public class ShoppingCart { ... }
```

But what if we only have one interceptor and we want the manager to ignore the value of `requiresNew` when binding interceptors? We can use the `@NonBinding` annotation:

```
@InterceptorBindingType
@Target({METHOD, TYPE})
@Retention(RUNTIME)
public @interface Secure {
    @NonBinding String[] rolesAllowed() default {};
}
```

7.5. Multiple interceptor binding annotations

Usually we use combinations of interceptor bindings types to bind multiple interceptors to a Web Bean. For example, the following declaration would be used to bind `TransactionInterceptor` and `SecurityInterceptor` to the same Web Bean:

```
@Secure(rolesAllowed="admin") @Transactional
public class ShoppingCart { ... }
```

However, in very complex cases, an interceptor itself may specify some combination of interceptor binding types:

```
@Transactional @Secure @Interceptor
public class TransactionalSecureInterceptor { ... }
```

Then this interceptor could be bound to the `checkout()` method using any one of the following combinations:

```
public class ShoppingCart {
    @Transactional @Secure public void checkout() { ... }
```

```
}
```

```
@Secure
public class ShoppingCart {
    @Transactional public void checkout() { ... }
}
```

```
@Transactional
public class ShoppingCart {
    @Secure public void checkout() { ... }
}
```

```
@Transactional @Secure
public class ShoppingCart {
    public void checkout() { ... }
}
```

7.6. Interceptor binding type inheritance

One limitation of the Java language support for annotations is the lack of annotation inheritance. Really, annotations should have reuse built in, to allow this kind of thing to work:

```
public @interface Action extends Transactional, Secure { ... }
```

Well, fortunately, Web Beans works around this missing feature of Java. We may annotate one interceptor binding type with other interceptor binding types. The interceptor bindings are transitive # any Web Bean with the first interceptor binding inherits the interceptor bindings declared as meta-annotations.

```
@Transactional @Secure
@InterceptorBindingType
@Target(TYPE)
@Retention(RUNTIME)
public @interface Action { ... }
```

Any Web Bean annotated `@Action` will be bound to both `TransactionInterceptor` and `SecurityInterceptor`. (And even `TransactionalSecureInterceptor`, if it exists.)

7.7. Use of `@Interceptors`

The `@Interceptors` annotation defined by the EJB specification is supported for both enterprise and simple Web Beans, for example:

```
@Interceptors({TransactionInterceptor.class, SecurityInterceptor.class})
public class ShoppingCart {
    public void checkout() { ... }
}
```

However, this approach suffers the following drawbacks:

- the interceptor implementation is hardcoded in business code,
- interceptors may not be easily disabled at deployment time, and
- the interceptor ordering is non-global # it is determined by the order in which interceptors are listed at the class level.

Therefore, we recommend the use of Web Beans-style interceptor bindings.

Decorators

Interceptors are a powerful way to capture and separate concerns which are *orthogonal* to the type system. Any interceptor is able to intercept invocations of any Java type. This makes them perfect for solving technical concerns such as transaction management and security. However, by nature, interceptors are unaware of the actual semantics of the events they intercept. Thus, interceptors aren't an appropriate tool for separating business-related concerns.

The reverse is true of *decorators*. A decorator intercepts invocations only for a certain Java interface, and is therefore aware of all the semantics attached to that interface. This makes decorators a perfect tool for modeling some kinds of business concerns. It also means that a decorator doesn't have the generality of an interceptor. Decorators aren't able to solve technical concerns that cut across many disparate types.

Suppose we have an interface that represents accounts:

```
public interface Account {
    public BigDecimal getBalance();
    public User getOwner();
    public void withdraw(BigDecimal amount);
    public void deposit(BigDecimal amount);
}
```

Several different Web Beans in our system implement the `Account` interface. However, we have a common legal requirement that, for any kind of account, large transactions must be recorded by the system in a special log. This is a perfect job for a decorator.

A decorator is a simple Web Bean that implements the type it decorates and is annotated `@Decorator`.

```
@Decorator
public abstract class LargeTransactionDecorator
    implements Account {

    @Decorates Account account;

    @PersistenceContext EntityManager em;

    public void withdraw(BigDecimal amount) {
        account.withdraw(amount);
        if ( amount.compareTo(LARGE_AMOUNT)>0 ) {
            em.persist( new LoggedWithdrawal(amount) );
        }
    }
}
```

```
}

public void deposit(BigDecimal amount);
    account.deposit(amount);
    if ( amount.compareTo(LARGE_AMOUNT)>0 ) {
        em.persist( new LoggedDeposit(amount) );
    }
}

}
```

Unlike other simple Web Beans, a decorator may be an abstract class. If there's nothing special the decorator needs to do for a particular method of the decorated interface, you don't need to implement that method.

8.1. Delegate attributes

All decorators have a *delegate attribute*. The type and binding types of the delegate attribute determine which Web Beans the decorator is bound to. The delegate attribute type must implement or extend all interfaces implemented by the decorator.

This delegate attribute specifies that the decorator is bound to all Web Beans that implement `Account`:

```
@Decorates Account account;
```

A delegate attribute may specify a binding annotation. Then the decorator will only be bound to Web Beans with the same binding.

```
@Decorates @Foreign Account account;
```

A decorator is bound to any Web Bean which:

- has the type of the delegate attribute as an API type, and
- has all binding types that are declared by the delegate attribute.

The decorator may invoke the delegate attribute, which has much the same effect as calling `InvocationContext.proceed()` from an interceptor.

8.2. Enabling decorators

We need to *enable* our decorator in `web-beans.xml`.

```
<Decorators>
  <myapp:LargeTransactionDecorator/>
</Decorators>
```

This declaration serves the same purpose for decorators that the `<Interceptors>` declaration serves for interceptors:

- it enables us to specify a total ordering for all decorators in our system, ensuring deterministic behavior, and
- it lets us enable or disable decorator classes at deployment time.

Interceptors for a method are called before decorators that apply to that method.

Events

The Web Beans event notification facility allows Web Beans to interact in a totally decoupled manner. Event *producers* raise events that are then delivered to event *observers* by the Web Bean manager. This basic schema might sound like the familiar observer/observable pattern, but there are a couple of twists:

- not only are event producers decoupled from observers; observers are completely decoupled from producers,
- observers can specify a combination of "selectors" to narrow the set of event notifications they will receive, and
- observers can be notified immediately, or can specify that delivery of the event should be delayed until the end of the current transaction

9.1. Event observers

An *observer method* is a method of a Web Bean with a parameter annotated `@Observes`.

```
public void onAnyDocumentEvent(@Observes Document document) { ... }
```

The annotated parameter is called the *event parameter*. The type of the event parameter is the observed *event type*. Observer methods may also specify "selectors", which are just instances of Web Beans binding types. When a binding type is used as an event selector, it is called an *event binding type*.

```
@BindingType  
@Target({PARAMETER, FIELD})  
@Retention(RUNTIME)  
public @interface Updated { ... }
```

We specify the event bindings of the observer method by annotating the event parameter:

```
public void afterDocumentUpdate(@Observes @Updated Document document) { ... }
```

An observer method need not specify any event bindings # in this case it is interested in *all* events of a particular type. If it does specify event bindings, it is only interested in events which also have those event bindings.

The observer method may have additional parameters, which are injected according to the usual Web Beans method parameter injection semantics:

```
public void afterDocumentUpdate(@Observes @Updated Document document, User user) { ... }
```

9.2. Event producers

The event producer may obtain an *event notifier* object by injection:

```
@Observable Event<Document> documentEvent
```

The `@Observable` annotation implicitly defines a Web Bean with scope `@Dependent` and deployment type `@Standard`, with an implementation provided by the Web Bean manager.

A producer raises events by calling the `fire()` method of the `Event` interface, passing an *event object*:

```
documentEvent.fire(document);
```

An event object may be an instance of any Java class that has no type variables or wildcard type parameters. The event will be delivered to every observer method that:

- has an event parameter to which the event object is assignable, and
- specifies no event bindings.

The Web Bean manager simply calls all the observer methods, passing the event object as the value of the event parameter. If any observer method throws an exception, the Web Bean manager stops calling observer methods, and the exception is rethrown by the `fire()` method.

To specify a "selector", the event producer may pass an instance of the event binding type to the `fire()` method:

```
documentEvent.fire( document, new AnnotationLiteral<Updated>(){} );
```

The helper class `AnnotationLiteral` makes it possible to instantiate binding types inline, since this is otherwise difficult to do in Java.

The event will be delivered to every observer method that:

- has an event parameter to which the event object is assignable, and

- does not specify any event binding *except* for the event bindings passed to `fire()`.

Alternatively, event bindings may be specified by annotating the event notifier injection point:

```
@Observable @Updated Event<Document> documentUpdatedEvent
```

Then every event fired via this instance of `Event` has the annotated event binding. The event will be delivered to every observer method that:

- has an event parameter to which the event object is assignable, and
- does not specify any event binding *except* for the event bindings passed to `fire()` or the annotated event bindings of the event notifier injection point.

9.3. Registering observers dynamically

It's often useful to register an event observer dynamically. The application may implement the `Observer` interface and register an instance with an event notifier by calling the `observe()` method.

```
documentEvent.observe( new Observer<Document>() { public void notify(Document doc) { ... }
});
```

Event binding types may be specified by the event notifier injection point or by passing event binding type instances to the `observe()` method:

```
documentEvent.observe( new Observer<Document>() { public void notify(Document doc) { ... } },
    new AnnotationLiteral<Updated>(){} );
```

9.4. Event bindings with members

An event binding type may have annotation members:

```
@BindingType
@Target({PARAMETER, FIELD})
@Retention(RUNTIME)
public @interface Role {
    RoleType value();
}
```

The member value is used to narrow the messages delivered to the observer:

```
public void adminLoggedIn(@Observes @Role(ADMIN) LoggedIn event) { ... }
```

Event binding type members may be specified statically by the event producer, via annotations at the event notifier injection point:

```
@Observable @Role(ADMIN) Event<LoggedIn> LoggedInEvent;}}
```

Alternatively, the value of the event binding type member may be determined dynamically by the event producer. We start by writing an abstract subclass of `AnnotationLiteral`:

```
abstract class RoleBinding  
    extends AnnotationLiteral<Role>  
    implements Role {}
```

The event producer passes an instance of this class to `fire()`:

```
documentEvent.fire( document, new RoleBinding() { public void value() { return user.getRole();  
    } } );
```

9.5. Multiple event bindings

Event binding types may be combined, for example:

```
@Observable @Blog Event<Document> blogEvent;  
...  
if (document.isBlog()) blogEvent.fire(document, new AnnotationLiteral<Updated>());
```

When this event occurs, all of the following observer methods will be notified:

```
public void afterBlogUpdate(@Observes @Updated @Blog Document document) { ... }
```

```
public void afterDocumentUpdate(@Observes @Updated Document document) { ... }
```

```
public void onAnyBlogEvent(@Observes @Blog Document document) { ... }
```

```
public void onAnyDocumentEvent(@Observes Document document) { ... }}
```

9.6. Transactional observers

Transactional observers receive their event notifications during the before or after completion phase of the transaction in which the event was raised. For example, the following observer method needs to refresh a query result set that is cached in the application context, but only when transactions that update the `Category` tree succeed:

```
public void refreshCategoryTree(@AfterTransactionSuccess @Observes CategoryUpdateEvent event) { ... }
```

There are three kinds of transactional observers:

- `@AfterTransactionSuccess` observers are called during the after completion phase of the transaction, but only if the transaction completes successfully
- `@AfterTransactionFailure` observers are called during the after completion phase of the transaction, but only if the transaction fails to complete successfully
- `@AfterTransactionCompletion` observers are called during the after completion phase of the transaction
- `@BeforeTransactionCompletion` observers are called during the before completion phase of the transaction

Transactional observers are very important in a stateful object model like Web Beans, because state is often held for longer than a single atomic transaction.

Imagine that we have cached a JPA query result set in the application scope:

```
@ApplicationScoped @Singleton
public class Catalog {

    @PersistenceContext EntityManager em;

    List<Product> products;

    @Produces @Catalog
```

```
List<Product> getCatalog() {  
    if (products==null) {  
        products = em.createQuery("select p from Product p where p.deleted = false")  
            .getResultList();  
    }  
    return products;  
}  
  
}
```

From time to time, a `Product` is created or deleted. When this occurs, we need to refresh the `Product` catalog. But we should wait until *after* the transaction completes successfully before performing this refresh!

The Web Bean that creates and deletes `Products` could raise events, for example:

```
@Stateless  
public class ProductManager {  
  
    @PersistenceContext EntityManager em;  
    @Observable Event<Product> productEvent;  
  
    public void delete(Product product) {  
        em.delete(product);  
        productEvent.fire(product, new AnnotationLiteral<Deleted>());  
    }  
  
    public void persist(Product product) {  
        em.persist(product);  
        productEvent.fire(product, new AnnotationLiteral<Created>());  
    }  
  
    ...  
  
}
```

And now `Catalog` can observe the events after successful completion of the transaction:

```
@ApplicationScoped @Singleton  
public class Catalog {  
  
    ...  
  
}
```

```
void addProduct(@AfterTransactionSuccess @Observes @Created Product product) {  
    products.add(product);  
}  
  
void addProduct(@AfterTransactionSuccess @Observes @Deleted Product product) {  
    products.remove(product);  
}  
  
}
```

Part III. Making the most of strong typing

The second major theme of Web Beans is *strong typing*. The information about the dependencies, interceptors and decorators of a Web Bean, and the information about event consumers for an event producer, is contained in typesafe Java constructs that may be validated by the compiler.

You don't see string-based identifiers in Web Beans code, not because the framework is hiding them from you using clever defaulting rules # so-called "configuration by convention" # but because there are simply no strings there to begin with!

The obvious benefit of this approach is that *any* IDE can provide autocompletion, validation and refactoring without the need for special tooling. But there is a second, less-immediately-obvious, benefit. It turns out that when you start thinking of identifying objects, events or interceptors via annotations instead of names, you have an opportunity to lift the semantic level of your code.

Web Beans encourages you develop annotations that model concepts, for example,

- `@Asynchronous`,
- `@Mock`,
- `@Secure` **or**
- `@Updated`,

instead of using compound names like

- `asyncPaymentProcessor`,
- `mockPaymentProcessor`,
- `SecurityInterceptor` **or**
- `DocumentUpdatedEvent`.

The annotations are reusable. They help describe common qualities of disparate parts of the system. They help us categorize and understand our code. They help us deal with common concerns in a common way. They make our code more literate and more understandable.

Web Beans *stereotypes* take this idea a step further. A stereotype models a common *role* in your application architecture. It encapsulates various properties of the role, including scope, interceptor bindings, deployment type, etc, into a single reusable package.

Even Web Beans XML metadata is strongly typed! There's no compiler for XML, so Web Beans takes advantage of XML schemas to validate the Java types and attributes that appear in XML.

This approach turns out to make the XML more literate, just like annotations made our Java code more literate.

We're now ready to meet some more advanced features of Web Beans. Bear in mind that these features exist to make our code both easier to validate and more understandable. Most of the time you don't ever really *need* to use these features, but if you use them wisely, you'll come to appreciate their power.

Stereotypes

According to the Web Beans specification:

In many systems, use of architectural patterns produces a set of recurring Web Bean roles. A stereotype allows a framework developer to identify such a role and declare some common metadata for Web Beans with that role in a central place.

A stereotype encapsulates any combination of:

- a default deployment type,
- a default scope type,
- a restriction upon the Web Bean scope,
- a requirement that the Web Bean implement or extend a certain type, and
- a set of interceptor binding annotations.

A stereotype may also specify that all Web Beans with the stereotype have defaulted Web Bean names.

A Web Bean may declare zero, one or multiple stereotypes.

A stereotype is a Java annotation type. This stereotype identifies action classes in some MVC framework:

```
@Retention(RUNTIME)
@Target(TYPE)
@Stereotype
public @interface Action {}
```

We use the stereotype by applying the annotation to a Web Bean.

```
@Action
public class LoginAction { ... }
```

10.1. Default scope and deployment type for a stereotype

A stereotype may specify a default scope and/or default deployment type for Web Beans with that stereotype. For example, if the deployment type `@WebTier` identifies Web Beans that should

only be deployed when the system executes as a web application, we might specify the following defaults for action classes:

```
@Retention(RUNTIME)
@Target(TYPE)
@RequestScoped
@WebTier
@Stereotype
public @interface Action {}
```

Of course, a particular action may still override these defaults if necessary:

```
@Dependent @Mock @Action
public class MockLoginAction { ... }
```

If we want to force all actions to a particular scope, we can do that too.

10.2. Restricting scope and type with a stereotype

Suppose that we wish to prevent actions from declaring certain scopes. Web Beans lets us explicitly specify the set of allowed scopes for Web Beans with a certain stereotype. For example:

```
@Retention(RUNTIME)
@Target(TYPE)
@RequestScoped
@WebTier
@Stereotype(supportedScopes=RequestScoped.class)
public @interface Action {}
```

If a particular action class attempts to specify a scope other than the Web Beans request scope, an exception will be thrown by the Web Bean manager at initialization time.

We can also force all Web Bean with a certain stereotype to implement an interface or extend a class:

```
@Retention(RUNTIME)
@Target(TYPE)
@RequestScoped
@WebTier
@Stereotype(requiredTypes=AbstractAction.class)
```

```
public @interface Action {}
```

If a particular action class does not extend the class `AbstractAction`, an exception will be thrown by the Web Bean manager at initialization time.

10.3. Interceptor bindings for stereotypes

A stereotype may specify a set of interceptor bindings to be inherited by all Web Beans with that stereotype.

```
@Retention(RUNTIME)
@Target(TYPE)
@RequestScoped
@Transactional(requiresNew=true)
@Secure
@WebTier
@Stereotype
public @interface Action {}
```

This helps us get technical concerns even further away from the business code!

10.4. Name defaulting with stereotypes

Finally, we can specify that all Web Beans with a certain stereotype have a Web Bean name, defaulted by the Web Bean manager. Actions are often referenced in JSP pages, so they're a perfect use case for this feature. All we need to do is add an empty `@Named` annotation:

```
@Retention(RUNTIME)
@Target(TYPE)
@RequestScoped
@Transactional(requiresNew=true)
@Secure
@Named
@WebTier
@Stereotype
public @interface Action {}
```

Now, `LoginAction` will have the name `loginAction`.

10.5. Standard stereotypes

We've already met two standard stereotypes defined by the Web Beans specification: `@Interceptor` and `@Decorator`.

Web Beans defines one further standard stereotype:

```
@Named
@RequestScoped
@Stereotype
@Target({TYPE, METHOD})
@Retention(RUNTIME)
public @interface Model {}
```

This stereotype is intended for use with JSF. Instead of using JSF managed beans, just annotate a Web Bean `@Model`, and use it directly in your JSF page.

Specialization

We've already seen how the Web Beans dependency injection model lets us *override* the implementation of an API at deployment time. For example, the following enterprise Web Bean provides an implementation of the API `PaymentProcessor` in production:

```
@CreditCard @Stateless
public class CreditCardPaymentProcessor
    implements PaymentProcessor {
    ...
}
```

But in our staging environment, we override that implementation of `PaymentProcessor` with a different Web Bean:

```
@CreditCard @Stateless @Staging
public class StagingCreditCardPaymentProcessor
    implements PaymentProcessor {
    ...
}
```

What we've tried to do with `StagingCreditCardPaymentProcessor` is to completely replace `AsyncPaymentProcessor` in a particular deployment of the system. In that deployment, the deployment type `@Staging` would have a higher priority than the default deployment type `@Production`, and therefore clients with the following injection point:

```
@CreditCard PaymentProcessor ccpp
```

Would receive an instance of `StagingCreditCardPaymentProcessor`.

Unfortunately, there are several traps we can easily fall into:

- the higher-priority Web Bean may not implement all the API types of the Web Bean that it attempts to override,
- the higher-priority Web Bean may not declare all the binding types of the Web Bean that it attempts to override,
- the higher-priority Web Bean might not have the same name as the Web Bean that it attempts to override, or

- the Web Bean that it attempts to override might declare a producer method, disposal method or observer method.

In each of these cases, the Web Bean that we tried to override could still be called at runtime. Therefore, overriding is somewhat prone to developer error.

Web Beans provides a special feature, called *specialization*, that helps the developer avoid these traps. Specialization looks a little esoteric at first, but it's easy to use in practice, and you'll really appreciate the extra security it provides.

11.1. Using specialization

Specialization is a feature that is specific to simple and enterprise Web Beans. To make use of specialization, the higher-priority Web Bean must:

- be a direct subclass of the Web Bean it overrides, and
- be a simple Web Bean if the Web Bean it overrides is a simple Web Bean or an enterprise Web Bean if the Web Bean it overrides is an enterprise Web Bean, and
- be annotated `@Specializes`.

```
@Stateless @Staging @Specializes
public class StagingCreditCardPaymentProcessor
    extends CreditCardPaymentProcessor {
    ...
}
```

We say that the higher-priority Web Bean *specializes* its superclass.

11.2. Advantages of specialization

When specialization is used:

- the binding types of the superclass are automatically inherited by the Web Bean annotated `@Specializes`, and
- the Web Bean name of the superclass is automatically inherited by the Web Bean annotated `@Specializes`, and
- producer methods, disposal methods and observer methods declared by the superclass are called upon an instance of the Web Bean annotated `@Specializes`.

In our example, the binding type `@CreditCard` of `CreditCardPaymentProcessor` is inherited by `StagingCreditCardPaymentProcessor`.

Furthermore, the Web Bean manager will validate that:

- all API types of the superclass are API types of the Web Bean annotated `@Specializes` (all local interfaces of the superclass enterprise bean are also local interfaces of the subclass),
- the deployment type of the Web Bean annotated `@Specializes` has a higher precedence than the deployment type of the superclass, and
- there is no other enabled Web Bean that also specializes the superclass.

If any of these conditions are violated, the Web Bean manager throws an exception at initialization time.

Therefore, we can be certain that the superclass will *never* be called in any deployment of the system where the Web Bean annotated `@Specializes` is deployed and enabled.

Defining Web Beans using XML

So far, we've seen plenty of examples of Web Beans declared using annotations. However, there are a couple of occasions when we can't use annotations to define the Web Bean:

- when the implementation class comes from some preexisting library, or
- when there should be multiple Web Beans with the same implementation class.

In either of these cases, Web Beans gives us two options:

- write a producer method, or
- declare the Web Bean using XML.

Many frameworks use XML to provide metadata relating to Java classes. However, Web Beans uses a very different approach to specifying the names of Java classes, fields or methods to most other frameworks. Instead of writing class and member names as the string values of XML elements and attributes, Web Beans lets you use the class or member name as the name of the XML element.

The advantage of this approach is that you can write an XML schema that prevents spelling errors in your XML document. It's even possible for a tool to generate the XML schema automatically from the compiled Java code. Or, an integrated development environment could perform the same validation without the need for the explicit intermediate generation step.

12.1. Declaring Web Bean classes

For each Java package, Web Beans defines a corresponding XML namespace. The namespace is formed by prepending `urn:java:` to the Java package name. For the package `com.mydomain.myapp`, the XML namespace is `urn:java:com.mydomain.myapp`.

Java types belonging to a package are referred to using an XML element in the namespace corresponding to the package. The name of the element is the name of the Java type. Fields and methods of the type are specified by child elements in the same namespace. If the type is an annotation, members are specified by attributes of the element.

For example, the element `<util:Date/>` in the following XML fragment refers to the class `java.util.Date`:

```
<WebBeans xmlns="urn:java:javax.webbeans"
  xmlns:util="urn:java:java.util">

  <util:Date/>

</WebBeans>
```

And this is all the code we need to declare that `Date` is a simple Web Bean! An instance of `Date` may now be injected by any other Web Bean:

```
@Current Date date
```

12.2. Declaring Web Bean metadata

We can declare the scope, deployment type and interceptor binding types using direct child elements of the Web Bean declaration:

```
<myapp:ShoppingCart>
  <SessionScoped/>
  <myfwk:Transactional requiresNew="true"/>
  <myfwk:Secure/>
</myapp:ShoppingCart>
```

We use exactly the same approach to specify names and binding type:

```
<util:Date>
  <Named>currentTime</Named>
</util:Date>

<util:Date>
  <SessionScoped/>
  <myapp:Login/>
  <Named>loginTime</Named>
</util:Date>

<util:Date>
  <ApplicationScoped/>
  <myapp:SystemStart/>
  <Named>systemStartTime</Named>
</util:Date>
```

Where `@Login` and `@SystemStart` are binding annotations types.

```
@Current Date currentTime;
@Login Date loginTime;
@SystemStart Date systemStartTime;
```

As usual, a Web Bean may support multiple binding types:

```
<myapp:AsynchronousChequePaymentProcessor>
  <myapp:PayByCheque/>
  <myapp:Asynchronous/>
</myapp:AsynchronousChequePaymentProcessor>
```

Interceptors and decorators are just simple Web Beans, so they may be declared just like any other simple Web Bean:

```
<myfwk:TransactionInterceptor>
  <Interceptor/>
  <myfwk:Transactional/>
</myfwk:TransactionInterceptor>
```

12.3. Declaring Web Bean members

TODO!

12.4. Declaring inline Web Beans

Web Beans lets us define a Web Bean at an injection point. For example:

```
<myapp:System>
  <ApplicationScoped/>
  <myapp:admin>
    <myapp:Name>
      <myapp:firstname>Gavin</myapp:firstname>
      <myapp:lastname>King</myapp:lastname>
      <myapp:email>gavin@hibernate.org</myapp:email>
    </myapp:Name>
  </myapp:admin>
</myapp:System>
```

The `<Name>` element declares a simple Web Bean of scope `@Dependent` and class `Name`, with a set of initial field values. This Web Bean has a special, container-generated binding and is therefore injectable only to the specific injection point at which it is declared.

This simple but powerful feature allows the Web Beans XML format to be used to specify whole graphs of Java objects. It's not quite a full databinding solution, but it's close!

12.5. Using a schema

If we want our XML document format to be authored by people who aren't Java developers, or who don't have access to our code, we need to provide a schema. There's nothing specific to Web Beans about writing or using the schema.

```
<WebBeans xmlns="urn:java:javax.webbeans"
  xmlns:myapp="urn:java:com.mydomain.myapp"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:java:javax.webbeans http://java.sun.com/jee/web-beans-1.0.xsd
    urn:java:com.mydomain.myapp http://mydomain.com/xsd/myapp-1.2.xsd">

  <myapp:System>
    ...
  </myapp:System>

</WebBeans>
```

Writing an XML schema is quite tedious. Therefore, the Web Beans RI project will provide a tool which automatically generates the XML schema from compiled Java code.

Part IV. Web Beans and the Java EE ecosystem

The third theme of Web Beans is *integration*. Web Beans was designed to work in concert with other technologies, helping the application developer fit the other technologies together. Web Beans is an open technology. It forms a part of the Java EE ecosystem, and is itself the foundation for a new ecosystem of portable extensions and integration with existing frameworks and technologies.

We've already seen how Web Beans helps integrate EJB and JSF, allowing EJBs to be bound directly to JSF pages. That's just the beginning. Web Beans offers the same potential to diverse other technologies, such as Business Process Management engines, other Web Frameworks, and third-party component models. The Java EE platform will never be able to standardize all the interesting technologies that are used in the world of Java application development, but Web Beans makes it easier to use the technologies which are not yet part of the platform seamlessly within the Java EE environment.

We're about to see how to take full advantage of the Java EE platform in an application that uses Web Beans. We'll also briefly meet a set of SPIs that are provided to support portable extensions to Web Beans. You might not ever need to use these SPIs directly, but it's nice to know they are there if you need them. Most importantly, you'll take advantage of them indirectly, every time you use a third-party extension.

Java EE integration

Web Beans is fully integrated into the Java EE environment. Web Beans have access to Java EE resources and JPA persistence contexts. They may be used in Unified EL expressions in JSF and JSP pages. They may even be injected into some objects, such as Servlets and Message-Driven Beans, which are not Web Beans.

13.1. Injecting Java EE resources into a Web Bean

All simple and enterprise Web Beans may take advantage of Java EE dependency injection using `@Resource`, `@EJB` and `@PersistenceContext`. We've already seen a couple of examples of this, though we didn't pay much attention at the time:

```
@Transactional @Interceptor
public class TransactionInterceptor {

    @Resource Transaction transaction;

    @AroundInvoke public Object manageTransaction(InvocationContext ctx) { ... }

}
```

```
@SessionScoped
public class Login {

    @Current Credentials credentials;
    @PersistenceContext EntityManager userDatabase;

    ...

}
```

The Java EE `@PostConstruct` and `@PreDestroy` callbacks are also supported for all simple and enterprise Web Beans. The `@PostConstruct` method is called after *all* injection has been performed.

There is one restriction to be aware of here: `@PersistenceContext(type=EXTENDED)` is not supported for simple Web Beans.

13.2. Calling a Web Bean from a Servlet

It's easy to use a Web Bean from a Servlet in Java EE 6. Simply inject the Web Bean using Web Beans field or initializer method injection.

```
public class Login extends HttpServlet {

    @Current Credentials credentials;
    @Current Login login;

    @Override
    public void service(HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {
        credentials.setUsername( request.getAttribute("username") );
        credentials.setPassword( request.getAttribute("password") );
        login.login();
        if ( login.isLoggedIn() ) {
            response.sendRedirect("/home.jsp");
        }
        else {
            response.sendRedirect("/loginError.jsp");
        }
    }
}
```

The Web Beans client proxy takes care of routing method invocations from the Servlet to the correct instances of `Credentials` and `Login` for the current request and HTTP session.

13.3. Calling a Web Bean from a Message-Driven Bean

Web Beans injection applies to all EJBs, even when they aren't under the control of the Web Bean manager (if they were obtained by direct JNDI lookup, or injection using `@EJB`, for example. In particular, you can use Web Beans injection in Message-Driven Beans, which are not considered Web Beans because you can't inject them.

You can even use Web Beans interceptor bindings for Message-Driven Beans.

```
@Transactional @MessageDriven
public class ProcessOrder implements MessageListener {

    @Current Inventory inventory;
    @PersistenceContext EntityManager em;
```

```

public void onMessage(Message message) {
    ...
}
}

```

Thus, receiving messages is super-easy in a Web Beans environment. But beware that there is no session or conversation context available when a message is delivered to a Message-Driven Bean. Only `@RequestScoped` and `@ApplicationScoped` Web Beans are available.

It's also easy to send messages using Web Beans.

13.4. JMS endpoints

Sending messages using JMS can be quite complex, because of the number of different objects you need to deal with. For queues we have `Queue`, `QueueConnectionFactory`, `QueueConnection`, `QueueSession` and `QueueSender`. For topics we have `Topic`, `TopicConnectionFactory`, `TopicConnection`, `TopicSession` and `TopicPublisher`. Each of these objects has its own lifecycle and threading model that we need to worry about.

Web Beans takes care of all this for us. All we need to do is declare the queue or topic in `web-beans.xml`, specifying an associated binding type and connection factory.

```

<Queue>
  <destination>java:comp/env/jms/OrderQueue</destination>
  <connectionFactory>java:comp/env/jms/QueueConnectionFactory</connectionFactory>
  <myapp:OrderProcessor/>
</Queue>

```

```

<Topic>
  <destination>java:comp/env/jms/StockPrices</destination>
  <connectionFactory>java:comp/env/jms/TopicConnectionFactory</connectionFactory>
  <myapp:StockPrices/>
</Topic>

```

Now we can just inject the `Queue`, `QueueConnection`, `QueueSession` or `QueueSender` for a queue, or the `Topic`, `TopicConnection`, `TopicSession` or `TopicPublisher` for a topic.

```

@OrderProcessor QueueSender orderSender;
@OrderProcessor QueueSession orderSession;

```

```
public void sendMessage() {  
    MapMessage msg = orderSession.createMapMessage();  
    ...  
    orderSender.send(msg);  
}
```

```
@StockPrices TopicPublisher pricePublisher;  
@StockPrices TopicSession priceSession;  
  
public void sendMessage(String price) {  
    pricePublisher.send( priceSession.createTextMessage(price) );  
}
```

The lifecycle of the injected JMS objects is completely controlled by the Web Bean manager.

13.5. Packaging and deployment

Web Beans doesn't define any special deployment archive. You can package Web Beans in JARs, EJB-JARs or WARs # any deployment location in the application classpath. However, each archive that contains Web Beans must include a file named `web-beans.xml` in the `META-INF` or `WEB-INF` directory. The file may be empty. Web Beans deployed in archives that do not have a `web-beans.xml` file will not be available for use in the application.

For Java SE execution, Web Beans may be deployed in any location in which EJBs may be deployed for execution by the embeddable EJB Lite container. Again, each location must contain a `web-beans.xml` file.

Extending Web Beans

Web Beans is intended to be a platform for frameworks, extensions and integration with other technologies. Therefore, Web Beans exposes a set of SPIs for the use of developers of portable extensions to Web Beans. For example, the following kinds of extensions were envisaged by the designers of Web Beans:

- integration with Business Process Management engines,
- integration with third-party frameworks such as Spring, Seam, GWT or Wicket, and
- new technology based upon the Web Beans programming model.

The nerve center for extending Web Beans is the `Manager` object.

14.1. The `Manager` object

The `Manager` interface lets us register and obtain Web Beans, interceptors, decorators, observers and contexts programmatically.

```
public interface Manager
{

    public <T> Set<Bean<T>> resolveByType(Class<T> type, Annotation... bindings);

    public <T> Set<Bean<T>> resolveByType(TypeLiteral<T> apiType,
        Annotation... bindings);

    public <T> T getInstanceByType(Class<T> type, Annotation... bindings);

    public <T> T getInstanceByType(TypeLiteral<T> type,
        Annotation... bindings);

    public Set<Bean<?>> resolveByName(String name);

    public Object getInstanceByName(String name);

    public <T> T getInstance(Been<T> bean);

    public void fireEvent(Object event, Annotation... bindings);

    public Context getContext(Class<? extends Annotation> scopeType);

    public Manager addContext(Context context);
```

```
public Manager addBean(Bean<?> bean);

public Manager addInterceptor(Interceptor interceptor);

public Manager addDecorator(Decorator decorator);

public <T> Manager addObserver(Observer<T> observer, Class<T> eventType,
    Annotation... bindings);

public <T> Manager addObserver(Observer<T> observer, TypeLiteral<T> eventType,
    Annotation... bindings);

public <T> Manager removeObserver(Observer<T> observer, Class<T> eventType,
    Annotation... bindings);

public <T> Manager removeObserver(Observer<T> observer,
    TypeLiteral<T> eventType, Annotation... bindings);

public <T> Set<Observer<T>> resolveObservers(T event, Annotation... bindings);

public List<Interceptor> resolveInterceptors(InterceptionType type,
    Annotation... interceptorBindings);

public List<Decorator> resolveDecorators(Set<Class<?>> types,
    Annotation... bindings);

}
```

We can obtain an instance of `Manager` via injection:

```
@Current Manager manager
```

14.2. The `Bean` class

Instances of the abstract class `Bean` represent Web Beans. There is an instance of `Bean` registered with the `Manager` object for every Web Bean in the application.

```
public abstract class Bean<T> {

    private final Manager manager;
```

```

protected Bean(Manager manager) {
    this.manager=manager;
}

protected Manager getManager() {
    return manager;
}

public abstract Set<Class> getTypes();
public abstract Set<Annotation> getBindingTypes();
public abstract Class<? extends Annotation> getScopeType();
public abstract Class<? extends Annotation> getDeploymentType();
public abstract String getName();

public abstract boolean isSerializable();
public abstract boolean isNullable();

public abstract T create();
public abstract void destroy(T instance);
}

```

It's possible to extend the `Bean` class and register instances by calling `Manager.addBean()` to provide support for new kinds of Web Beans, beyond those defined by the Web Beans specification (simple and enterprise Web Beans, producer methods and JMS endpoints). For example, we could use the `Bean` class to allow objects managed by another framework to be injected into Web Beans.

There are two subclasses of `Bean` defined by the Web Beans specification: `Interceptor` and `Decorator`.

14.3. The Context interface

The `Context` interface supports addition of new scopes to Web Beans, or extension of the built-in scopes to new environments.

```

public interface Context {

    public Class<? extends Annotation> getScopeType();

    public <T> T get(Bean<T> bean, boolean create);

    boolean isActive();
}

```

```
}
```

For example, we might implement `Context` to add a business process scope to Web Beans, or to add support for the conversation scope to an application that uses Wicket.

Next steps

Because Web Beans is so new, there's not yet a lot of information available online.

Of course, the Web Beans specification is the best source of more information about Web Beans. The spec is about 100 pages long, only twice the length of this article, and almost as readable. But, of course, it covers many details that we've skipped over. The spec is available from <http://jcp.org/en/jsr/detail?id=299>.

The Web Beans Reference implementation is being developed at <http://seamframework.org/WebBeans>. The RI development team and the Web Beans spec lead blog at <http://in.relation.to>. This article is substantially based upon a series of blog entries published there.

Appendix A. Integrating the Web Beans RI into other environments

Currently the Web Beans RI only runs in JBoss AS 5; integrating the RI into other EE environments (for example another application server like Glassfish), into a servlet container (like Tomcat), or with an Embedded EJB3.1 implementation is fairly easy. In this Appendix we will briefly discuss the steps needed.



Note

It should be possible to run Web Beans in an SE environment, but you'll to do more work, adding your own contexts and lifecycle. The Web Beans RI currently doesn't expose lifecycle extension points, so you would have to code directly against Web Beans RI classes.

A.1. The Web Beans RI SPI

The Web Beans SPI is located in `webbeans-ri-spi` module, and packaged as `webbeans-ri-spi.jar`. Some SPIs are optional, if you need to override the default behavior, others are required.

You can specify the implementation of an SPI either as a system property, or in a properties file `META-INF/web-beans-ri.properties`. All property names are the fully qualified class name of the implemented interface; all property values are the fully qualified class name of the implementation class.

All interfaces in the SPI support the decorator pattern and provide a `Forwarding` class.

A.1.1. Web Bean Discovery

```
public interface WebBeanDiscovery {  
    /**  
     * Gets list of all classes in classpath archives with web-beans.xml files  
     *  
     * @return An iterable over the classes  
     */  
    public Iterable<Class<?>> discoverWebBeanClasses();  
  
    /**  
     * Gets a list of all web-beans.xml files in the app classpath  
     *  
     * @return An iterable over the web-beans.xml files  
     */  
}
```

```
public Iterable<URL> discoverWebBeansXml();  
  
}
```

The discovery of Web Bean classes and `web-bean.xml` files is self-explanatory (the algorithm is described in Section 11.1 of the JSR-299 specification, and isn't repeated here).

The Web Beans RI can be told to load your implementation of `WebBeanDiscovery` using the property `org.jboss.webbeans.bootstrap.spi.WebBeanDiscovery` with the fully qualified class name as the value. For example:

```
org.jboss.webbeans.bootstrap.spi.WebBeanDiscovery=org.jboss.webbeans.integration.jbossas.WebBeanDiscovery
```

If the Web Beans RI is being used in a servlet container, it expects a constructor of the form:

```
public WebBeanDiscoveryImpl(ServletContext servletContext) {}
```

The servlet context can be used to allow your implementation of `WebBeanDiscovery` to interact with the container.

A.1.2. EJB Discovery

The Web Beans RI also delegates EJB3 bean discovery to the container so that it doesn't have to scan for EJB3 annotations or parse `ejb-jar.xml`. For each EJB in the application an `EJBDescriptor` should be discovered:

```
public interface EjbDiscovery  
{  
    public static final String PROPERTY_NAME = EjbDiscovery.class.getName();  
  
    /**  
     * Gets a descriptor for each EJB in the application  
     *  
     * @return The bean class to descriptor map  
     */  
    public Iterable<EjbDescriptor<?>> discoverEjbs();  
}
```

```
public interface EjbDescriptor<T> {
```

```
/**
 * Gets the EJB type
 *
 * @return The EJB Bean class
 */
public Class<T> getType();

/**
 * Gets the local business interfaces of the EJB
 *
 * @return An iterator over the local business interfaces
 */
public Iterable<BusinessInterfaceDescriptor<?>> getLocalBusinessInterfaces();

/**
 * Gets the remote business interfaces of the EJB
 *
 * @return An iterator over the remote business interfaces
 */
public Iterable<BusinessInterfaceDescriptor<?>> getRemoteBusinessInterfaces();

/**
 * Get the remove methods of the EJB
 *
 * @return An iterator over the remove methods
 */
public Iterable<Method> getRemoveMethods();

/**
 * Indicates if the bean is stateless
 *
 * @return True if stateless, false otherwise
 */
public boolean isStateless();

/**
 * Indicates if the bean is a EJB 3.1 Singleton
 *
 * @return True if the bean is a singleton, false otherwise
 */
public boolean isSingleton();

/**
```

```
* Indicates if the EJB is stateful
*
* @return True if the bean is stateful, false otherwise
*/
public boolean isStateful();

/**
* Indicates if the EJB is and MDB
*
* @return True if the bean is an MDB, false otherwise
*/
public boolean isMessageDriven();

/**
* Gets the EJB name
*
* @return The name
*/
public String getEjbName();

}
```

The `EjbDescriptor` is fairly self-explanatory, and should return the relevant metadata as defined in the EJB specification. In addition to these two interfaces, there is `BusinessInterfaceDescriptor` which represents a local business interface (encapsulating the interface class and jndi name used to look up an instance of the EJB).

The Web Beans RI can be told to load your implementation of `EjbDiscovery` using the property `org.jboss.webbeans.bootstrap.spi.EjbDiscovery` with the fully qualified class name as the value. For example:

```
org.jboss.webbeans.bootstrap.spi.EjbDiscovery=org.jboss.webbeans.integration.jbossas.EjbDiscoveryImpl
```

If the Web Beans RI is being used in a servlet container, it expects a constructor of the form:

```
public EjbDiscoveryImpl(ServletContext servletContext) {}
```

The servlet context can be used to allow your implementation of `EjbDiscovery` to interact with the container.

A.1.3. JNDI

The Web Beans RI implements JNDI binding and lookup according to standards, however you may want to alter the binding and lookup (for example in an environment where JNDI isn't available). To do this, implement `org.jboss.webbeans.spi.resources.NamingContext`:

```
public interface NamingContext extends Serializable {

    /**
     * Typed JNDI lookup
     *
     * @param <T> The type
     * @param name The JNDI name
     * @param expectedType The expected type
     * @return The object
     */
    public <T> T lookup(String name, Class<? extends T> expectedType);

    /**
     * Binds an item to JNDI
     *
     * @param name The key to bind under
     * @param value The item to bind
     */
    public void bind(String name, Object value);

}
```

and tell the RI to use it:

```
org.jboss.webbeans.resources.spi.NamingContext=com.acme.MyNamingContext
```

If the Web Beans RI is being used in a servlet container, it expects a constructor of the form:

```
public MyNamingContext(ServletContext servletContext) {}
```

The servlet context can be used to allow your implementation of `NamingContext` to interact with the container.

A.1.4. Resource loading

The Web Beans RI needs to load classes and resources from the classpath at various times. By default, they are loaded from the same classloader that was used to load the RI, however this may not be correct for some environments. If this is case, you can implement `org.jboss.webbeans.spi.ResourceLoader`:

```
public interface ResourceLoader {

    /**
     * Creates a class from a given FQCN
     *
     * @param name The name of the clas
     * @return The class
     */
    public Class<?> classForName(String name);

    /**
     * Gets a resource as a URL by name
     *
     * @param name The name of the resource
     * @return An URL to the resource
     */
    public URL getResource(String name);

    /**
     * Gets resources as URLs by name
     *
     * @param name The name of the resource
     * @return An iterable reference to the URLs
     */
    public Iterable<URL> getResources(String name);

}
```

and tell the RI to use it:

```
org.jboss.webbeans.resources.spi.ResourceLoader=com.acme.ResourceLoader
```

If the Web Beans RI is being used in a servlet container, it expects a constructor of the form:


```
public MyResourceLoader(ServletContext servletContext) {}
```

The servlet context can be used to allow your implementation of `ResourceLoader` to interact with the container.

A.2. The contract with the container

There are a number of requirements that the Web Beans RI places on the container for correct functioning that fall outside implementation of APIs

Classloader isolation

If you are integrating the Web Beans RI into an environment that supports deployment of multiple applications, you must enable, automatically, or through user configuration, classloader isolation for each Web Beans application.

Servlet listener

If you are integrating the Web Beans into a Servlet environment you must register `org.jboss.webbeans.servlet.WebBeansListener` as a Servlet listener, either automatically, or through user configuration, for each Web Beans application which uses Servlet.

Session Bean Interceptor

If you are integrating the Web Beans into an EJB environment you must register `org.jboss.webbeans.ejb.SessionBeanInterceptor` as a EJB interceptor for all EJBs in the application, either automatically, or through user configuration, for each Web Beans application which uses enterprise beans.

The `webbeans-ri.jar`

If you are integrating the Web Beans into an environment that supports deployment of applications, you must insert the `webbeans-ri.jar` into the applications isolated classloader. It cannot be loaded from a shared classloader.

