

Weld - JSR-299 Reference Implementation

JSR-299: The new Java standard for dependency injection and contextual lifecycle management

Gavin King

Pete Muir

Dan Allen

David Allen

Italian Translation: Nicola Benaglia, Francesco Milesi

Spanish Translation: Gladys Guerrero

Korean Translation: Eun-Ju Ki,

Traditional Chinese Translation: Terry Chuang

Simplified Chinese Translation: Sean Wu

A note about naming and nomenclature	vii
I. Beans	1
1. Introduction	3
1.1. What is a bean?	3
1.2. Getting our feet wet	3
2. More about beans	7
2.1. The anatomy of a bean	7
2.1.1. Bean types, qualifiers and dependency injection	8
2.1.2. Scope	10
2.1.3. EL name	10
2.1.4. Alternatives	11
2.1.5. Interceptor binding types	11
2.2. What kinds of classes are beans?	12
2.2.1. Managed beans	12
2.2.2. Session beans	13
2.2.3. Producer methods	14
2.2.4. Producer fields	16
3. JSF web application example	17
4. Dependency injection and programmatic lookup	21
4.1. Injection points	21
4.2. What gets injected	22
4.3. Qualifier annotations	23
4.4. The built-in qualifiers @Default and @Any	24
4.5. Qualifiers with members	25
4.6. Multiple qualifiers	25
4.7. Alternatives	26
4.8. Fixing unsatisfied and ambiguous dependencies	26
4.9. Client proxies	27
4.10. Obtaining a contextual instance by programmatic lookup	28
4.11. The InjectionPoint object	29
5. Scopes and contexts	33
5.1. Scope types	33
5.2. Built-in scopes	33
5.3. The conversation scope	34
5.3.1. Conversation demarcation	35
5.3.2. Conversation propagation	35
5.3.3. Conversation timeout	36
5.4. The singleton pseudo-scope	36
5.5. The dependent pseudo-scope	37
5.6. The @New qualifier	37
II. Getting Start with Weld, the CDI Reference Implementation	39
6. Getting started with Weld	41
6.1. Prerequisites	41
6.2. Deploying to JBoss AS	41
6.3. Deploying to GlassFish	43
6.4. Deploying to Apache Tomcat	44
6.4.1. Deploying with Ant	45
6.4.2. Deploying with Maven	46
6.5. Deploying to Jetty	47
7. Diving into the Weld examples	49
7.1. The numberguess example in depth	49
7.1.1. The numberguess example in Apache Tomcat or Jetty	54
7.2. The numberguess example for Java SE with Swing	54

7.2.1. Creating the Eclipse project	54
7.2.2. Running the example from Eclipse	54
7.2.3. Running the example from the command line	57
7.2.4. Understanding the code	57
7.3. The translator example in depth	62
III. Loose coupling with strong typing	67
8. Producer methods	69
8.1. Scope of a producer method	70
8.2. Injection into producer methods	70
8.3. Use of @New with producer methods	71
8.4. Disposer methods	71
9. Interceptors	73
9.1. Interceptor bindings	73
9.2. Implementing interceptors	74
9.3. Enabling interceptors	74
9.4. Interceptor bindings with members	75
9.5. Multiple interceptor binding annotations	76
9.6. Interceptor binding type inheritance	77
9.7. Use of @Interceptors	77
10. Decorators	79
10.1. Delegate object	80
10.2. Enabling decorators	81
11. Events	83
11.1. Event payload	83
11.2. Event observers	83
11.3. Event producers	84
11.4. Conditional observer methods	85
11.5. Event qualifiers with members	85
11.6. Multiple event qualifiers	86
11.7. Transactional observers	86
12. Stereotypes	89
12.1. Default scope for a stereotype	89
12.2. Interceptor bindings for stereotypes	90
12.3. Name defaulting with stereotypes	90
12.4. Alternative stereotypes	90
12.5. Stereotype stacking	91
12.6. Built-in stereotypes	91
13. Specialization, inheritance and alternatives	93
13.1. Using alternative stereotypes	93
13.2. A minor problem with alternatives	94
13.3. Using specialization	95
14. Java EE component environment resources	97
14.1. Defining a resource	97
14.2. Typesafe resource injection	98
IV. CDI and the Java EE ecosystem	101
15. Java EE integration	103
15.1. Built-in beans	103
15.2. Injecting Java EE resources into a bean	103
15.3. Calling a bean from a servlet	104
15.4. Calling a bean from a message-driven bean	104
15.5. JMS endpoints	105
15.6. Packaging and deployment	106
16. Portable extensions	107

16.1. Creating an Extension	107
16.2. Container lifecycle events	108
16.3. The BeanManager object	109
16.4. The InjectionTarget interface	109
16.5. The Bean interface	110
16.6. Registering a Bean	111
16.7. Wrapping an AnnotatedType	112
16.8. Wrapping an InjectionTarget	115
16.9. The Context interface	117
17. Next steps	119
V. Weld Reference Guide	121
18. Application servers and environments supported by Weld	123
18.1. Using Weld with JBoss AS	123
18.2. GlassFish	123
18.3. Servlet containers (such as Tomcat or Jetty)	123
18.3.1. Tomcat	124
18.3.2. Jetty	124
18.4. Java SE	125
18.4.1. CDI SE Module	126
18.4.2. Bootstrapping CDI SE	126
18.4.3. Thread Context	128
18.4.4. Setting the Classpath	128
19. Context Management	131
19.1. Managing the built in contexts	131
20. Configuration	135
20.1. Excluding classes from scanning and deployment	135
20.2. Parallel bootstrap configuration	136
A. Integrating Weld into other environments	139
A.1. The Weld SPI	139
A.1.1. Deployment structure	139
A.1.2. EJB descriptors	141
A.1.3. EE resource injection and resolution services	141
A.1.4. EJB services	141
A.1.5. JPA services	142
A.1.6. Transaction Services	142
A.1.7. Resource Services	142
A.1.8. Injection Services	142
A.1.9. Security Services	142
A.1.10. Bean Validation Services	142
A.1.11. Identifying the BDA being addressed	142
A.1.12. The bean store	143
A.1.13. The application context	143
A.1.14. Initialization and shutdown	143
A.1.15. Resource loading	143
A.2. The contract with the container	144

A note about naming and nomenclature

Shortly before the final draft of JSR-299 was submitted, the specification changed its name from "Web Beans" to "Java Contexts and Dependency Injection for the Java EE platform", abbreviated CDI. For a brief period after the renaming, the reference implementation adopted the name "Web Beans". However, this ended up causing more confusion than it solved and Red Hat decided to change the name of the reference implementation to "Weld". You may still find other documentation, blogs, forum posts, etc. that use the old nomenclature. Please update any references you can. The naming game is over.

You'll also find that some of the functionality that once existed in the specification is now missing, such as defining beans in XML. These features will be available as portable extensions for CDI in the Weld project, and perhaps other implementations.

Note that this reference guide was started while changes were still being made to the specification. We've done our best to update it for accuracy. If you discover a conflict between what is written in this guide and the specification, the specification is the authority—assume it is correct. If you believe you have found an error in the specification, please report it to the JSR-299 EG.

Part I. Beans

The [JSR-299](http://jcp.org/en/jsr/detail?id=299) [http://jcp.org/en/jsr/detail?id=299] specification (CDI) defines a set of complementary services that help improve the structure of application code. CDI layers an enhanced lifecycle and interaction model over existing Java component types, including managed beans and Enterprise Java Beans. The CDI services provide:

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

The CDI services are a core aspect of the Java EE platform and include full support for Java EE modularity and the Java EE component architecture. But the specification does not limit the use of CDI to the Java EE environment. In the Java SE environment, the services might be provided by a standalone CDI implementation like Weld (see [Section 18.4.1, “CDI SE Module”](#)), or even by a container that also implements the subset of EJB defined for embedded usage by the EJB 3.1 specification. CDI is especially useful in the context of web application development, but the problems it solves are general development concerns and it is therefore applicable to a wide variety of application.

An object bound to a lifecycle context is called a bean. CDI includes built-in support for several different kinds of bean, including the following Java EE component types:

- managed beans, and
- EJB session beans.

Both managed beans and EJB session beans may inject other beans. But some other objects, which are not themselves beans in the sense used here, may also have beans injected via CDI. In the Java EE platform, the following kinds of component may have beans injected:

- message-driven beans,
- interceptors,
- servlets, servlet filters and servlet event listeners,
- JAX-WS service endpoints and handlers, and
- JSP tag handlers and tag library event listeners.

CDI relieves the user of an unfamiliar API of the need to answer the following questions:

- What is the lifecycle of this object?
 - How many simultaneous clients can it have?
 - Is it multithreaded?
 - How do I get access to it from a client?
-

- Do I need to explicitly destroy it?
- Where should I keep the reference to it when I'm not currently using it?
- How can I define an alternative implementation, so that the implementation can vary at deployment time?
- How should I go about sharing this object between other objects?

CDI is more than a framework. It's a whole, rich programming model. The *theme* of CDI is *loose-coupling with strong typing*. Let's study what that phrase means.

A bean specifies only the type and semantics of other beans it depends upon. It need not be aware of the actual lifecycle, concrete implementation, threading model or other clients of any bean it interacts with. Even better, the concrete implementation, lifecycle and threading model of a bean may vary according to the deployment scenario, without affecting any client. This loose-coupling makes your code easier to maintain.

Events, interceptors and decorators enhance the loose-coupling 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.

What's even more powerful (and comforting) is that CDI provides all these facilities in a *typesafe* way. CDI never relies on string-based identifiers to determine how collaborating objects fit together. Instead, CDI uses the typing information that is already available in the Java object model, augmented using a new programming pattern, called *qualifier annotations*, to wire together beans, their dependencies, their interceptors and decorators, and their event consumers. Usage of XML descriptors is minimized to truly deployment-specific information.

But CDI isn't a restrictive programming model. It doesn't tell you how you should to structure your application into layers, how you should handle persistence, or what web framework you have to use. You'll have to decide those kinds of things for yourself.

CDI even provides a comprehensive SPI, allowing other kinds of object defined by future Java EE specifications or by third-party frameworks to be cleanly integrated with CDI, take advantage of the CDI services, and interact with any other kind of bean.

CDI was influenced by a number of existing Java frameworks, including Seam, Guice and Spring. However, CDI 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. But it couldn't have been any of these without inspiration from the frameworks mentioned and *lots* of collaboration and hard work by the JSR-299 Expert Group (EG).

Finally, CDI is a [Java Community Process](http://jcp.org) [http://jcp.org] (JCP) standard. Java EE 6 requires that all compliant application servers provide support for JSR-299 (even in the web profile).

Introduction

So you're keen to get started writing your first bean? Or perhaps you're skeptical, wondering what kinds of hoops the CDI specification will make you jump through! The good news is that you've probably already written and used hundreds, perhaps thousands of beans. CDI just makes it easier to actually use them to build an application!

1.1. What is a bean?

A bean is exactly what you think it is. Only now, it has a true identity in the container environment.

Prior to Java EE 6, there was no clear definition of the term "bean" in the Java EE platform. Of course, we've been calling Java classes used in web and enterprise applications "beans" for years. There were even a couple of different kinds of things called "beans" in EE specifications, including EJB beans and JSF managed beans. Meanwhile, other third-party frameworks such as Spring and Seam introduced their own ideas of what it meant to be a "bean". What we've been missing is a common definition.

Java EE 6 finally lays down that common definition in the Managed Beans specification. Managed Beans are defined as container-managed objects with minimal programming restrictions, otherwise known by the acronym POJO (Plain Old Java Object). They support a small set of basic services, such as resource injection, lifecycle callbacks and interceptors. Companion specifications, such as EJB and CDI, build on this basic model. But, *at last*, there's a uniform concept of a bean and a lightweight component model that's aligned across the Java EE platform.

With very few exceptions, almost every concrete Java class that has a constructor with no parameters (or a constructor designated with the annotation `@Inject`) is a bean. This includes every `JavaBean` and every EJB session bean. If you've already got some `JavaBeans` or session beans lying around, they're already beans—you won't need any additional special metadata. There's just little one thing you need to do before you can start injecting them into stuff: you need to put them in an archive (a jar, or a Java EE module such as a war or EJB jar) that contains a special marker file: `META-INF/beans.xml`.

The `JavaBeans` and EJBs you've been writing every day, up until now, have not been able to take advantage of the new services defined by the CDI specification. But you'll be able to use every one of them with CDI—allowing the container to create and destroy instances of your beans and associate them with a designated context, injecting them into other beans, using them in EL expressions, specializing them with qualifier annotations, even adding interceptors and decorators to them—without modifying your existing code. At most, you'll need to add some annotations.

Now let's see how to create your first bean that actually uses CDI.

1.2. Getting our feet wet

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 EJB local interface:

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

Unfortunately, we don't have a class that translates whole text documents. So let's write a bean for this job:

```
public class TextTranslator {  
    private SentenceParser sentenceParser;  
    private Translator sentenceTranslator;  
  
    @Inject  
    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();  
    }  
}
```

But wait! `TextTranslator` does not have a constructor with no parameters! Is it still a bean? If you remember, a class that does not have a constructor with no parameters can still be a bean if it has a constructor annotated `@Inject`.

As you've guessed, the `@Inject` annotation has something to do with dependency injection! `@Inject` may be applied to a constructor or method of a bean, and tells the container to call that constructor or method when instantiating the bean. The container will inject other beans into the parameters of the constructor or method.

We may obtain an instance of `TextTranslator` by injecting it into a constructor, method or field of a bean, or a field or method of a Java EE component class such as a servlet. The container chooses the object to be injected based on the type of the injection point, not the name of the field, method or parameter.

Let's create a UI controller bean that uses field injection to obtain an instance of the `TextTranslator`, translating the text entered by a user:

```
@Named @RequestScoped  
public class TranslateController {  
    @Inject TextTranslator textTranslator;  
  
    private String inputText;
```

①

```

private String translation;

// JSF action method, perhaps
public void translate() {
    translation = textTranslator.translate(inputText);
}

public String getInputText() {
    return inputText;
}

public void setInputText(String text) {
    this.inputText = text;
}

public String getTranslation() {
    return translation;
}
}

```

1 Field injection of TextTranslator instance



Tip

Notice the controller bean is request-scoped and named. Since this combination is so common in web applications, there's a built-in annotation for it in CDI that we could have used as a shorthand. When the (stereotype) annotation `@Model` is declared on a class, it creates a request-scoped and named bean.

Alternatively, we may obtain an instance of `TextTranslator` programmatically from an injected instance of `Instance`, parameterized with the bean type:

```

@Inject Instance<TextTranslator> textTranslatorInstance;
...
public void translate() {
    textTranslatorInstance.get().translate(inputText);
}

```

Notice that it isn't necessary to create a getter or setter method to inject one bean into another. CDI can access an injected field directly (even if it's private!), which sometimes helps eliminate some wasteful code. The name of the field is arbitrary. It's the field's type that determines what is injected.

At system initialization time, the container must validate that exactly one bean exists which satisfies each injection point. In our example, if no implementation of `Translator` is available—if the `SentenceTranslator` EJB was not deployed—the container would inform us of an *unsatisfied dependency*. If more than one implementation of `Translator` were available, the container would inform us of the *ambiguous dependency*.

Before we get too deep in the details, let's pause and examine a bean's anatomy. What aspects of the bean are significant, and what gives it its identity? Instead of just giving examples of beans, we're going to define what *makes* something a bean.

More about beans

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

Let's back up a second. What does it really mean to be *contextual*? Since 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 bean see the bean in different states. The client-visible state depends upon which instance of the 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 bean determines:

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

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

Clients (for example, other beans) executing in the same context will see the same instance of the bean. But clients in a different context may see a different instance (depending on the relationship between the contexts).

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

- they interact via well-defined public APIs
- their lifecycles are completely decoupled

We can replace one bean with another different bean that implements the same interface and has a different lifecycle (a different scope) without affecting the other bean implementation. In fact, CDI defines a simple facility for overriding bean implementations at deployment time, as we will see in [Section 4.7, "Alternatives"](#).

Note that not all clients of a bean are beans themselves. Other objects such as servlets or message-driven beans—which are by nature not injectable, contextual objects—may also obtain references to beans by injection.

2.1. The anatomy of a bean

Enough hand-waving. More formally, the anatomy of a bean, according to the spec:

A bean comprises the following attributes:

- A (nonempty) set of bean types
- A (nonempty) set of qualifiers

- A scope
- Optionally, a bean EL name
- A set of interceptor bindings
- A bean implementation

Furthermore, a bean may or may not be an alternative.

Let's see what all this new terminology means.

2.1.1. Bean types, qualifiers and dependency injection

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

- a bean type, together with
- a set of qualifiers.

A bean type is a user-defined class or interface; a type that is client-visible. If the bean is an EJB session bean, the bean type is the `@Local` interface or bean-class local view. A bean may have multiple bean types. For example, the following bean has four bean types:

```
public class BookShop
    extends Business
    implements Shop<Book> {
    ...
}
```

The bean types are `BookShop`, `Business` and `Shop<Book>`, as well as the implicit type `java.lang.Object`. (Notice that a parameterized type is a legal bean type).

Meanwhile, this session bean has only the local interfaces `BookShop`, `Auditable` and `java.lang.Object` as bean types, since the bean class, `BookShopBean` is not a client-visible type.

```
@Stateful
public class BookShopBean
    extends Business
    implements BookShop, Auditable {
    ...
}
```



Note

The bean types of a session bean include local interfaces and the bean class local view (if any). EJB remote interfaces are not considered bean types of a session bean. You can't inject an EJB using its remote interface unless you define a *resource*, which we'll meet in [Chapter 14, Java EE component environment resources](#).

Bean types may be restricted to an explicit set by annotating the bean with the `@Typed` annotation and listing the classes that should be bean types. For instance, the bean types of this bean have been restricted to `Shop<Book>`, together with `java.lang.Object`:

```
@Typed(Shop.class)
public class BookShop
    extends Business
    implements Shop<Book> {
    ...
}
```

Sometimes, a bean type alone does not provide enough information for the container to know which bean to inject. For instance, suppose we have two implementations of the `PaymentProcessor` interface: `CreditCardPaymentProcessor` and `DebitPaymentProcessor`. Injecting a field of type `PaymentProcessor` introduces an ambiguous condition. In these cases, the client must specify some additional quality of the implementation it is interested in. We model this kind of "quality" using a qualifier.

A qualifier is a user-defined annotation that is itself annotated `@Qualifier`. A qualifier annotation is an extension of the type system. It lets us disambiguate a type without having to fall back to string-based names. Here's an example of a qualifier annotation:

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

You may not be used to seeing the definition of an annotation. In fact, this might be the first time you've encountered one. With CDI, annotation definitions will become a familiar artifact as you'll be creating them from time to time.



Note

Pay attention to the names of the built-in annotations in CDI and EJB. You'll notice that they are often adjectives. We encourage you to follow this convention when creating your custom annotations, since they serve to describe the behaviors and roles of the class.

Now that we have defined a qualifier annotation, we can use it to disambiguate an injection point. The following injection point has the bean type `PaymentProcessor` and qualifier `@CreditCard`:

```
@Inject @CreditCard PaymentProcessor paymentProcessor
```

For each injection point, the container searches for a bean which satisfies the contract, one which has the bean type and all the qualifiers. If it finds exactly one matching bean, it injects an instance of that bean. If it doesn't, it reports an error to the user.

How do we specify that qualifiers of a bean? By annotating the bean class, of course! The following bean has the qualifier `@CreditCard` and implements the bean type `PaymentProcessor`. Therefore, it satisfies our qualified injection point:

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



Note

If a bean or an injection point does not explicitly specify a qualifier, it has the default qualifier, `@Default`.

That's not quite the end of the story. CDI also defines a simple *resolution rule* that helps the container decide what to do if there is more than one bean that satisfies a particular contract. We'll get into the details in [Chapter 4, Dependency injection and programmatic lookup](#).

2.1.2. Scope

The *scope* of a bean defines the lifecycle and visibility of its instances. The CDI context model is extensible, accommodating arbitrary scopes. However, certain important scopes are built into the specification, and provided by the container. Each scope is represented by an annotation type.

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

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

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



Note

Keep in mind that once a bean is bound to a context, it remains in that context until the context is destroyed. There is no way to manually remove a bean from a context. If you don't want the bean to sit in the session indefinitely, consider using another scope with a shorted lifespan, such as the request or conversation scope.

If a scope is not explicitly specified, then the bean belongs to a special scope called the *dependent pseudo-scope*. Beans with this scope live to serve the object into which they were injected, which means their lifecycle is bound to the lifecycle of that object.

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

2.1.3. EL name

If you want to reference a bean in non-Java code that supports Unified EL expressions, for example, in a JSP or JSF page, you must assign the bean an *EL name*.

The EL name is specified using the `@Named` annotation, as shown here:

```
public @SessionScoped @Named("cart")
```

```
class ShoppingCart implements Serializable { ... }
```

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

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



Note

The `@Named` annotation is not what makes the class a bean. Most classes in a bean archive are already recognized as beans. The `@Named` annotation just makes it possible to reference the bean from the EL, most commonly from a JSF view.

We can let CDI choose a name for us by leaving off the value of the `@Named` annotation:

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

The name defaults to the unqualified class name, decapitalized; in this case, `shoppingCart`.

2.1.4. Alternatives

We've already seen how qualifiers let us choose between multiple implementations of an interface at development time. But sometimes we have an interface (or other bean type) whose implementation varies depending upon the deployment environment. For example, we may want to use a mock implementation in a testing environment. An *alternative* may be declared by annotating the bean class with the `@Alternative` annotation.

```
public @Alternative
class MockPaymentProcessor extends PaymentProcessorImpl { ... }
```

We normally annotate a bean `@Alternative` only when there is some other implementation of an interface it implements (or of any of its bean types). We can choose between alternatives at deployment time by *selecting* an alternative in the CDI deployment descriptor `META-INF/beans.xml` of the jar or Java EE module that uses it. Different modules can specify that they use different alternatives.

We cover alternatives in more detail in [Section 4.7, "Alternatives"](#).

2.1.5. Interceptor binding types

You might be familiar with the use of interceptors in EJB 3.0. In Java EE 6, this functionality has been generalized to work with other managed beans. That's right, you no longer have to make your bean an EJB just to intercept its methods. Holler. So what does CDI have to offer above and beyond that? Well, quite a lot actually. Let's cover some background.

The way that interceptors were defined in Java EE 5 was counter-intuitive. You were required to specify the *implementation* of the interceptor directly on the *implementation* of the EJB, either in the `@Interceptors`

annotation or in the XML descriptor. You might as well just put the interceptor code *in* the implementation! Second, the order in which the interceptors are applied is taken from the order in which they are declared in the annotation or the XML descriptor. Perhaps this isn't so bad if you're applying the interceptors to a single bean. But, if you are applying them repeatedly, then there's a good chance that you'll inadvertently define a different order for different beans. Now that's a problem.

CDI provides a new approach to binding interceptors to beans that introduces a level of indirection (and thus control). We must define an *interceptor binding type* to describe the behavior implemented by the interceptor.

An interceptor binding type is a user-defined annotation that is itself annotated `@InterceptorBinding`. It lets us bind interceptor classes to bean classes with no direct dependency between the two classes.

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

The interceptor that implements transaction management declares this annotation:

```
public @Transactional @Interceptor
class TransactionInterceptor { ... }
```

We can apply the interceptor to a bean by annotating the bean class with the same interceptor binding type:

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

Notice that `ShoppingCart` and `TransactionInterceptor` don't know anything about each other.

Interceptors are deployment-specific. (We don't need a `TransactionInterceptor` in our unit tests!) By default, an interceptor is disabled. We can enable an interceptor using the CDI deployment descriptor `META-INF/beans.xml` of the jar or Java EE module. This is also where we specify the interceptor ordering.

We'll discuss interceptors, and their cousins, decorators, in [Chapter 9, Interceptors](#) and [Chapter 10, Decorators](#).

2.2. What kinds of classes are beans?

We've already seen two types of beans: JavaBeans and EJB session beans. Is that the whole story? Actually, it's just the beginning. Let's explore the various kinds of beans that CDI implementations must support out-of-the-box.

2.2.1. Managed beans

A managed bean is a Java class. The basic lifecycle and semantics of a managed bean are defined by the Managed Beans specification. You can explicitly declare a managed bean by annotating the bean class `@ManagedBean`, but in CDI you don't need to. According to the specification, the CDI container treats any class that satisfies the following conditions as a managed bean:

- It is not a non-static inner class.

- It is a concrete class, or is annotated `@Decorator`.
- It is not annotated with an EJB component-defining annotation or declared as an EJB bean class in `ejb-jar.xml`.
- It does not implement `javax.enterprise.inject.spi.Extension`.
- It has an appropriate constructor—either:
 - the class has a constructor with no parameters, or
 - the class declares a constructor annotated `@Inject`.



Note

According to this definition, JPA entities are technically managed beans. However, entities have their own special lifecycle, state and identity model and are usually instantiated by JPA or using `new`. Therefore we don't recommend directly injecting an entity class. We especially recommend against assigning a scope other than `@Dependent` to an entity class, since JPA is not able to persist injected CDI proxies.

The unrestricted set of bean types for a managed bean contains the bean class, every superclass and all interfaces it implements directly or indirectly.

If a managed bean has a public field, it must have the default scope `@Dependent`.

Managed beans support the `@PostConstruct` and `@PreDestroy` lifecycle callbacks.

Session beans are also, technically, managed beans. However, since they have their own special lifecycle and take advantage of additional enterprise services, the CDI specification considers them to be a different kind of bean.

2.2.2. Session beans

Session beans belong to the EJB specification. They have a special lifecycle, state management and concurrency model that is different to other managed beans and non-managed Java objects. But session beans participate in CDI just like any other bean. You can inject one session bean into another session bean, a managed bean into a session bean, a session bean into a managed bean, have a managed bean observe an event raised by a session bean, and so on.



Note

Message-driven and entity beans are by nature non-contextual objects and may not be injected into other objects. However, message-driven beans can take advantage of some CDI functionality, such as dependency injection, interceptors and decorators. In fact, CDI will perform injection into any session or message-driven bean, even those which are not contextual instances.

The unrestricted set of bean types for a session bean contains all local interfaces of the bean and their superinterfaces. If the session bean has a bean class local view, the unrestricted set of bean types contains the bean class and all superclasses. In addition, `java.lang.Object` is a bean type of every session bean. But remote interfaces are *not* included in the set of bean types.

There's no reason to explicitly declare the scope of a stateless session bean or singleton session bean. The EJB container controls the lifecycle of these beans, according to the semantics of the `@Stateless` or `@Singleton` declaration. On the other hand, a stateful session bean may have any scope.

Stateful session beans may define a *remove method*, annotated `@Remove`, that is used by the application to indicate that an instance should be destroyed. However, for a contextual instance of the bean—an instance under the control of CDI—this method may only be called by the application if the bean has scope `@Dependent`. For beans with other scopes, the application must let the container destroy the bean.

So, when should we use a session bean instead of a plain managed bean? 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 or web service invocation, or
- timers and asynchronous methods,

When we don't need any of these things, an ordinary managed bean will serve just fine.

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

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.

The point we're trying to make is: use a session bean when you need the services it provides, not just because you want to use dependency injection, lifecycle management, or interceptors. Java EE 6 provides a graduated programming model. It's usually easy to start with an ordinary managed bean, and later turn it into an EJB just by adding one of the following annotations: `@Stateless`, `@Stateful` or `@Singleton`.

On the other hand, don't be scared to use session beans just because you've heard your friends say they're "heavyweight". It's nothing more than superstition to think that something is "heavier" just because it's hosted natively within the Java EE container, instead of by a proprietary bean container or dependency injection framework that runs as an additional layer of obfuscation. And as a general principle, you should be skeptical of folks who use vaguely defined terminology like "heavyweight".

2.2.3. Producer methods

Not everything that needs to be injected can be boiled down to a bean class instantiated by the container using `new`. There are plenty of cases where we need additional control. What if we need to decide at runtime which implementation of a type to instantiate and inject? What if we need to inject an object that is obtained by querying a service or transactional resource, for example by executing a JPA query?

A *producer method* is a method that acts as a source of bean instances. The method declaration itself describes the bean and the container invokes the method to obtain an instance of the bean when no instance exists in the specified context. A producer method lets the application take full control of the bean instantiation process.

A producer method is declared by annotating a method of a bean class with the `@Produces` annotation.

```
@ApplicationScoped
```

```
public class RandomNumberGenerator {

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

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

}
```

We can't write a bean class that is itself a random number. But we can certainly write a method that returns a random number. By making the method a producer method, we allow the return value of the method—in this case an `Integer`—to be injected. We can even specify a qualifier—in this case `@Random`, a scope—which in this case defaults to `@Dependent`, and an EL name—which in this case defaults to `randomNumber` according to the JavaBeans property name convention. Now we can get a random number anywhere:

```
@Inject @Random int randomNumber;
```

Even in a Unified EL expression:

```
<p>Your raffle number is #{randomNumber}</p>
```

A producer method must be a non-abstract method of a managed bean class or session bean class. A producer method may be either static or non-static. If the bean is a session bean, the producer method must be either a business method of the EJB or a static method of the bean class.

The bean types of a producer method depend upon the method return type:

- If the return type is an interface, the unrestricted set of bean types contains the return type, all interfaces it extends directly or indirectly and `java.lang.Object`.
- If a return type is primitive or is a Java array type, the unrestricted set of bean types contains exactly two types: the method return type and `java.lang.Object`.
- If the return type is a class, the unrestricted set of bean types contains the return type, every superclass and all interfaces it implements directly or indirectly.



Note

Producer methods and fields may have a primitive bean type. For the purpose of resolving dependencies, primitive types are considered to be identical to their corresponding wrapper types in `java.lang`.

If the producer method has method parameters, the container will look for a bean that satisfies the type and qualifiers of each parameter and pass it to the method automatically—another form of dependency injection.

```
@Produces Set<Roles> getRoles(User user) {
    return user.getRoles();
}
```

```
}
```

We'll talk much more about producer methods in [Chapter 8, Producer methods](#).

2.2.4. Producer fields

A *producer field* is a simpler alternative to a producer method. A producer field is declared by annotating a field of a bean class with the `@Produces` annotation—the same annotation used for producer methods.

```
public class Shop {  
    @Produces PaymentProcessor paymentProcessor = ....;  
    @Produces @Catalog List<Product> products = ....;  
}
```

The rules for determining the bean types of a producer field parallel the rules for producer methods.

A producer field is really just a shortcut that lets us avoid writing a useless getter method. However, in addition to convenience, producer fields serve a specific purpose as an adaptor for Java EE component environment injection, but to learn more about that, you'll have to wait until [Chapter 14, Java EE component environment resources](#). Because we can't wait to get to work on some examples.

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 request-scoped bean to hold the username and password entered during login, with constraints defined using annotations from the Bean Validation specification:

```
@Named @RequestScoped
public class Credentials {
    private String username;
    private String password;

    @NotNull @Length(min=3, max=25)
    public String getUsername() { return username; }
    public void setUsername(String username) { this.username = username; }

    @NotNull @Length(min=6, max=20)
    public String getPassword() { return password; }
    public void setPassword(String password) { this.password = password; }
}
```

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

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

Users are represented by a JPA entity:

```
@Entity
public class User {
    private @NotNull @Length(min=3, max=25) @Id String username;
    private @NotNull @Length(min=6, max=20) String password;

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

(Note that we're also going to need a `persistence.xml` file to configure the JPA persistence unit containing `User`.)

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

```
@SessionScoped @Named
public class Login implements Serializable {

    @Inject Credentials credentials;
    @Inject @UserDatabase 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);
        }
        else {
            // perhaps add code here to report a failed login
        }
    }

    public void logout() {
        user = null;
    }

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

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

`@LoggedIn` and `@UserDatabase` are custom qualifier annotations:

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

```
@Qualifier
@Retention(RUNTIME)
@Target({METHOD, PARAMETER, FIELD})
public @interface UserDatabase {}
```

We need an adaptor bean to expose our typesafe `EntityManager`:

```
class UserDatabaseProducer {
    @Produces @UserDatabase @PersistenceContext
    static EntityManager userDatabase;
}
```

Now `DocumentEditor`, or any other bean, can easily inject the current user:

```
public class DocumentEditor {
    @Inject Document document;
    @Inject @LoggedIn User currentUser;
    @Inject @DocumentDatabase EntityManager docDatabase;

    public void save() {
        document.setCreatedBy(currentUser);
        docDatabase.persist(document);
    }
}
```

Or we can reference the current user in a JSF view:

```
<h:panelGroup rendered="#{login.loggedIn}">
    signed in as #{currentUser.username}
</h:panelGroup>
```

Hopefully, this example gave you a taste of the CDI programming model. In the next chapter, we'll explore dependency injection in greater depth.

Dependency injection and programmatic lookup

One of the most significant features of CDI—certainly the most recognized—is dependency injection; excuse me, *typesafe* dependency injection.

4.1. Injection points

The `@Inject` annotation lets us define an injection point that is injected during bean instantiation. Injection can occur via three different mechanisms.

Bean constructor parameter injection:

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

A bean can only have one injectable constructor.

Initializer method parameter injection:

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



Note

A bean can have multiple initializer methods. If the bean is a session bean, the initializer method is not required to be a business method of the session bean.

And direct field injection:

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



Note

Getter and setter methods are not required for field injection to work (unlike with JSF managed beans).

Dependency injection always occurs when the bean instance is first instantiated by the container. Simplifying just a little, things happen in this order:

- First, the container calls the bean constructor (the default constructor or the one annotated `@Inject`), to obtain an instance of the bean.
- Next, the container initializes the values of all injected fields of the bean.
- Next, the container calls all initializer methods of bean (the call order is not portable, don't rely on it).
- Finally, the `@PostConstruct` method, if any, is called.

(The only complication is that the container might call initializer methods declared by a superclass before initializing injected fields declared by a subclass.)



Note

One major advantage of constructor injection is that it allows the bean to be immutable.

CDI also supports parameter injection for some other methods that are invoked by the container. For instance, parameter injection is supported for producer methods:

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

This is a case where the `@Inject` annotation *is not* required at the injection point. The same is true for observer methods (which we'll meet in [Chapter 11, Events](#)) and disposer methods.

4.2. What gets injected

The CDI specification defines a procedure, called *typesafe resolution*, that the container follows when identifying the 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 container will inform the developer immediately if a bean's dependencies cannot be satisfied.

The purpose of this algorithm is to allow multiple beans to implement the same bean type and either:

- allow the client to select which implementation it requires using a *qualifier* or
- allow the application deployer to select which implementation is appropriate for a particular deployment, without changes to the client, by enabling or disabling an *alternative*, or
- allow the beans to be isolated into separate modules.

Obviously, if you have exactly one bean of a given type, and an injection point with that same type, then bean A is going to go into slot A. That's the simplest possible scenario. When you first start your application, you'll likely have lots of those.

But then, things start to get complicated. Let's explore how the container determines which bean to inject in more advanced cases. We'll start by taking a closer look at qualifiers.

4.3. Qualifier annotations

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

```
@Synchronous
public class SynchronousPaymentProcessor implements PaymentProcessor {
    public void process(Payment payment) { ... }
}
```

```
@Asynchronous
public class AsynchronousPaymentProcessor implements PaymentProcessor {
    public void process(Payment payment) { ... }
}
```

Where `@Synchronous` and `@Asynchronous` are qualifier annotations:

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

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

A client bean developer uses the qualifier annotation to specify exactly which bean should be injected.

Using field injection:

```
@Inject @Synchronous PaymentProcessor syncPaymentProcessor;
```

```
@Inject @Asynchronous PaymentProcessor asyncPaymentProcessor;
```

Using initializer method injection:

```
@Inject
public void setPaymentProcessors(@Synchronous PaymentProcessor syncPaymentProcessor,
                                @Asynchronous PaymentProcessor asyncPaymentProcessor) {
    this.syncPaymentProcessor = syncPaymentProcessor;
    this.asyncPaymentProcessor = asyncPaymentProcessor;
}
```

Using constructor injection:

```
@Inject
public Checkout(@Synchronous PaymentProcessor syncPaymentProcessor,
                @Asynchronous PaymentProcessor asyncPaymentProcessor) {
    this.syncPaymentProcessor = syncPaymentProcessor;
    this.asyncPaymentProcessor = asyncPaymentProcessor;
}
```

Qualifier annotations can also qualify method arguments of producer, disposer and observer methods. Combining qualified arguments with producer methods is a good way to have an implementation of a bean type selected at runtime based on the state of the system:

```
@Produces
PaymentProcessor getPaymentProcessor(@Synchronous PaymentProcessor syncPaymentProcessor,
                                     @Asynchronous PaymentProcessor asyncPaymentProcessor) {
    return isSynchronous() ? syncPaymentProcessor : asyncPaymentProcessor;
}
```

If an injected field or a parameter of a bean constructor or initializer method is not explicitly annotated with a qualifier, the default qualifier, `@Default`, is assumed.

Now, you may be thinking, *"What's the different between using a qualifier and just specifying the exact implementation class you want?"* It's important to understand that a qualifier is like an extension of the interface. It does not create a direct dependency to any particular implementation. There may be multiple alternative implementations of `@Asynchronous PaymentProcessor`!

4.4. The built-in qualifiers `@Default` and `@Any`

Whenever a bean or injection point does not explicitly declare a qualifier, the container assumes the qualifier `@Default`. From time to time, you'll need to declare an injection point without specifying a qualifier. There's a qualifier for that too. All beans have the qualifier `@Any`. Therefore, by explicitly specifying `@Any` at an injection point, you suppress the default qualifier, without otherwise restricting the beans that are eligible for injection.

This is especially useful if you want to iterate over all beans with a certain bean type. For example:

```
@Inject
```



```
void initServices(@Any Instance<Service> services) {
    for (Service service: services) {
        service.init();
    }
}
```

4.5. Qualifiers with members

Java annotations can have members. We can use annotation members to further discriminate a qualifier. This prevents a potential explosion of new annotations. For example, instead of creating several qualifiers representing different payment methods, we could aggregate them into a single annotation with a member:

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

Then we select one of the possible member values when applying the qualifier:

```
private @Inject @PayBy(CHECK) PaymentProcessor checkPayment;
```

We can force the container to ignore a member of a qualifier type by annotating the member `@Nonbinding`.

```
@Qualifier
@Retention(RUNTIME)
@Target({METHOD, FIELD, PARAMETER, TYPE})
public @interface PayBy {
    PaymentMethod value();
    @Nonbinding String comment() default "";
}
```

4.6. Multiple qualifiers

An injection point may specify multiple qualifiers:

```
@Inject @Synchronous @Reliable PaymentProcessor syncPaymentProcessor;
```

Then only a bean which has *both* qualifier annotations would be eligible for injection.

```
@Synchronous @Reliable
public class SynchronousReliablePaymentProcessor implements PaymentProcessor {
```

```
public void process(Payment payment) { ... }
}
```

4.7. Alternatives

Alternatives are beans whose implementation is specific to a particular client module or deployment scenario. This alternative defines a mock implementation of both `@Synchronous PaymentProcessor` and `@Asynchronous PaymentProcessor`, all in one:

```
@Alternative @Synchronous @Asynchronous
public class MockPaymentProcessor implements PaymentProcessor {
    public void process(Payment payment) { ... }
}
```

By default, `@Alternative` beans are disabled. We need to *enable* an alternative in the `beans.xml` descriptor of a bean archive to make it available for instantiation and injection. This activation only applies to the beans in that archive.

```
<beans
  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/beans_1_0.xsd">
  <alternatives>
    <class>org.mycompany.mock.MockPaymentProcessor</class>
  </alternatives>
</beans>
```

When an ambiguous dependency exists at an injection point, the container attempts to resolve the ambiguity by looking for an enabled alternative among the beans that could be injected. If there is exactly one enabled alternative, that's the bean that will be injected.

4.8. Fixing unsatisfied and ambiguous dependencies

The typesafe resolution algorithm fails when, after considering the qualifier annotations on all beans that implement the bean type of an injection point and filtering out disabled beans (`@Alternative` beans which are not explicitly enabled), the container is unable to identify exactly one bean to inject. The container will abort deployment, informing us of the unsatisfied or ambiguous dependency.

During the course of your development, you're going to encounter this situation. Let's learn how to resolve it.

To fix an *unsatisfied dependency*, either:

- create a bean which implements the bean type and has all the qualifier types of the injection point,
- make sure that the bean you already have is in the classpath of the module with the injection point, or
- explicitly enable an `@Alternative` bean that implements the bean type and has the appropriate qualifier types, using `beans.xml`.

To fix an *ambiguous dependency*, either:

- introduce a qualifier to distinguish between the two implementations of the bean type,
- disable one of the beans by annotating it `@Alternative`,
- move one of the implementations to a module that is not in the classpath of the module with the injection point, or
- disable one of two `@Alternative` beans that are trying to occupy the same space, using `beans.xml`.

See [this](#) [FAQ](#) [<http://sfwk.org/Documentation/HowDoAResolveAnAmbiguousResolutionExceptionBetweenAProducerMethodAndARawType>] for step-by-step instructions for how to resolve an ambiguous resolution exception between a raw bean type and a producer method that returns the same bean type.

Just remember: "There can be only one."

On the other hand, if you really do have an optional or multivalued injection point, you should change the type of your injection point to `Instance`, as we'll see in [Section 4.10](#), "[Obtaining a contextual instance by programmatic lookup](#)".

Now there's one more issue you need to be aware of when using the dependency injection service.

4.9. Client proxies

Clients of an injected bean do not usually hold a direct reference to a bean instance, unless the bean is a dependent object (`scope@Dependent`).

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

Now imagine that a bean bound to the session scope holds a direct reference to a 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 bean instance should not be serialized along with the session scoped bean! It can get that reference any time. No need to hoard it!

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

Unfortunately, due to limitations of the Java language, some Java types cannot be proxied by the container. If an injection point declared with one of these types resolves to a bean with any scope other than `@Dependent`, the container will abort deployment, informing us of the problem.

The following Java types cannot be proxied by the container:

- classes which don't have a non-private constructor with no parameters, and
- classes which are declared `final` or have a `final` method,
- arrays and primitive types.

It's usually very easy to fix an unproxyable dependency problem. If an injection point of type `X` results in an unproxyable dependency, simply:

- add a constructor with no parameters to `X`,
- change the type of the injection point to `Instance<X>`,
- introduce an interface `Y`, implemented by the injected bean, and change the type of the injection point to `Y`, or

- if all else fails, change the scope of the injected bean to `@Dependent`.



Note

A future release of Weld will likely support a non-standard workaround for this limitation, using non-portable JVM APIs:

- Sun, IcedTea, Mac: `Unsafe.allocateInstance()` (The most efficient)
- IBM, JRockit: `ReflectionFactory.newConstructorForSerialization()`

But we didn't get around to implementing this yet.

4.10. Obtaining a contextual instance by programmatic lookup

In certain situations, injection is not the most convenient way to obtain a contextual reference. For example, it may not be used when:

- the bean type or qualifiers vary dynamically at runtime, or
- depending upon the deployment, there may be no bean which satisfies the type and qualifiers, or
- we would like to iterate over all beans of a certain type.

In these situations, the application may obtain an instance of the interface `Instance`, parameterized for the bean type, by injection:

```
@Inject Instance<PaymentProcessor> paymentProcessorSource;
```

The `get()` method of `Instance` produces a contextual instance of the bean.

```
PaymentProcessor p = paymentProcessorSource.get();
```

Qualifiers can be specified in one of two ways:

- by annotating the `Instance` injection point, or
- by passing qualifiers to the `select()` of `Event`.

Specifying the qualifiers at the injection point is much, much easier:

```
@Inject @Asynchronous Instance<PaymentProcessor> paymentProcessorSource;
```

Now, the `PaymentProcessor` returned by `get()` will have the qualifier `@Asynchronous`.

Alternatively, we can specify the qualifier dynamically. First, we add the `@Any` qualifier to the injection point, to suppress the default qualifier. (All beans have the qualifier `@Any`.)

```
@Inject @Any Instance<PaymentProcessor> paymentProcessorSource;
```

Next, we need to obtain an instance of our qualifier type. Since annotations are interfaces, we can't just write `new Asynchronous()`. It's also quite tedious to create a concrete implementation of an annotation type from scratch. Instead, CDI lets us obtain a qualifier instance by subclassing the helper class `AnnotationLiteral`.

```
class AsynchronousQualifier
extends AnnotationLiteral<Asynchronous> implements Asynchronous {}
```

In some cases, we can use an anonymous class:

```
PaymentProcessor p = paymentProcessorSource
    .select(new AnnotationLiteral<Asynchronous>() {});
```

However, we can't use an anonymous class to implement a qualifier type with members.

Now, finally, we can pass the qualifier to the `select()` method of `Instance`.

```
Annotation qualifier = synchronously ?
    new SynchronousQualifier() : new AsynchronousQualifier();
PaymentProcessor p = anyPaymentProcessor.select(qualifier).get().process(payment);
```

4.11. The `InjectionPoint` object

There are certain kinds of dependent objects (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 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:

```
@Inject Logger log;
```

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

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

We would use this qualifier 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(InjectionPoint ip) {
        ServletRequest request = (ServletRequest) FacesContext.getCurrentInstance().getExternalContext().getRequest
        return request.getParameter(ip.getAnnotated().getAnnotation(HttpParam.class).value());
    }

}
```

Note that acquiring of the request in this example is JSF-centric. For a more generic solution you could write your own producer for the request and have it injected as a method parameter.

Note also that the `value()` member of the `HttpParam` annotation is ignored by the container since it is annotated `@Nonbinding`.

The container provides a built-in bean that implements the `InjectionPoint` interface:

```
public interface InjectionPoint {  
    public Type getType();  
    public Set<Annotation> getQualifiers();  
    public Bean<?> getBean();  
    public Member getMember();  
    public Annotated getAnnotated();  
    public boolean isDelegate();  
    public boolean isTransient();  
}
```


Scopes and contexts

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

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

For example, if we have a session-scoped bean, `CurrentUser`, all 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.



Note

JPA entities aren't a great fit for this model. Entities have their whole own lifecycle and identity model which just doesn't map naturally to the model used in CDI. Therefore, we recommend against treating entities as CDI beans. You're certainly going to run into problems if you try to give an entity a scope other than the default scope `@Dependent`. The client proxy will get in the way if you try to pass an injected instance to the JPA `EntityManager`.

5.1. Scope types

CDI features an *extensible context model*. It's possible to define new scopes by creating a new scope type annotation:

```
@ScopeType
@Retention(RUNTIME)
@Target({TYPE, METHOD})
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. You can expect an implementation of the business scope, for instance, in a future version of Seam.

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

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

Usually, you'll use one of CDI's built-in scopes.

5.2. Built-in scopes

CDI defines four built-in scopes:

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

For a web application that uses CDI:

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

A CDI extension can implement support for the conversation scope in other web frameworks.

The request and application scopes are also active:

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

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

Managed beans with scope `@SessionScoped` or `@ConversationScoped` must be serializable, since the container passivates the HTTP session from time to time.

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 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 (browsers tend to share domain cookies, and hence the session cookie, between tabs, so this is not the case for the session scope).

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. 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

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

```
@Inject 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 bean controls the conversation with which it is associated:

```
@ConversationScoped @Stateful
public class OrderBuilder {
    private Order order;
    private @Inject 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 bean is able to control its own lifecycle through use of the `Conversation` API. But some other 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) or redirect. 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 CDI 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 EL bean name `javax.enterprise.context.conversation`.

Therefore, the following link propagates the conversation:

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

It's probably better to use one of the link components in JSF 2:

```
<h:link outcome="/addProduct.xhtml" value="Add Product">
  <f:param name="cid" value="#{javax.enterprise.context.conversation.id}" />
</h:link>
```



Tip

The conversation context propagates across redirects, making it very easy to implement the common POST-then-redirect pattern, without resort to fragile constructs such as a "flash" object. The container automatically adds the conversation id to the redirect URL as a request parameter.

5.3.3. Conversation timeout

The container is permitted to destroy a conversation and all state held in its context at any time in order to conserve resources. A CDI implementation will normally do this on the basis of some kind of timeout—though this is not required by the specification. The timeout is the period of inactivity before the conversation is destroyed (as opposed to the amount of time the conversation is active).

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

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

5.4. The singleton pseudo-scope

In addition to the four built-in scopes, CDI also supports two *pseudo-scopes*. The first is the *singleton pseudo-scope*, which we specify using the annotation `@Singleton`.



Note

Unlike the other scopes, which belong to the package `javax.enterprise.context`, the `@Singleton` annotation is defined in the package `javax.inject`.

You can guess what "singleton" means here. It means a bean that is instantiated once. Unfortunately, there's a little problem with this pseudo-scope. Beans with scope `@Singleton` don't have a proxy object. Clients hold a direct reference to the singleton instance. So we need to consider the case of a client that can be serialized, for example,

any bean with scope `@SessionScoped` or `@ConversationScoped`, any dependent object of a bean with scope `@SessionScoped` or `@ConversationScoped`, or any stateful session bean.

Now, if the singleton instance is a simple, immutable, serializable object like a string, a number or a date, we probably don't mind too much if it gets duplicated via serialization. However, that makes it no stop being a true singleton, and we may as well have just declared it with the default scope.

There are several ways to ensure that the singleton bean remains a singleton when its client gets serialized:

- have the singleton bean implement `writeResolve()` and `readReplace()` (as defined by the Java serialization specification),
- make sure the client keeps only a transient reference to the singleton bean, or
- give the client a reference of type `Instance<X>` where `X` is the bean type of the singleton bean.

A fourth, better solution is to instead use `@ApplicationScoped`, allowing the container to proxy the bean, and take care of serialization problems automatically.

5.5. The dependent pseudo-scope

Finally, CDI features the so-called *dependent pseudo-scope*. This is the default scope for a bean which does not explicitly declare a scope type.

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

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

An instance of a dependent bean is never shared between different clients or different injection points. It is strictly a *dependent object* of some other object. It is instantiated when the object it belongs to is created, and destroyed when the object it belongs to is destroyed.

If a Unified EL expression refers to a dependent bean by EL name, an instance of the bean is instantiated every time the expression is evaluated. The instance is not reused during any other expression evaluation.



Note

If you need to access a bean directly by EL name in a JSF page, you probably need to give it a scope other than `@Dependent`. Otherwise, any value that gets set to the bean by a JSF input will be lost immediately. That's why CDI features the `@Model` stereotype; it lets you give a bean a name, and set its scope to `@RequestScoped` in one stroke. If you need to access a bean that really *has* to have the scope `@Dependent` from a JSF page, inject it into a different bean, and expose it to EL via a getter method.

Beans with scope `@Dependent` don't need a proxy object. The client holds a direct reference to its instance.

CDI makes it easy to obtain a dependent instance of a bean, even if the bean is already declared as a bean with some other scope type.

5.6. The `@New` qualifier

The built-in qualifier `@New` allows us to obtain a dependent object of a specified class.

```
@Inject @New Calculator calculator;
```

The class must be a valid managed bean or session bean, but need not be an enabled bean.

This works 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 {  
    @Inject Calculator calculator;  
    @Inject @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 [Chapter 8, Producer methods](#).

Part II. Getting Start with Weld, the CDI Reference Implementation

Weld, the JSR-299 Reference Implementation (RI), is being developed as part of the [Seam project](http://seamframework.org/Weld) [http://seamframework.org/Weld]. You can download the latest community release of Weld from the [download page](http://seamframework.org/Download) [http://seamframework.org/Download]. Information about the Weld source code repository and instructions about how to obtain and build the source can be found on the same page.

Weld provides a complete SPI allowing Java EE containers such as JBoss AS and GlassFish to use Weld as their built-in CDI implementation. Weld also runs in servlet engines like Tomcat and Jetty, or even in a plain Java SE environment.

Weld comes with an extensive library of examples, which are a great starting point from which to learn CDI.

Getting started with Weld

Weld comes with a number of examples. We recommend you start with `examples/jsf/numberguess` and `examples/jsf/translator`. Numberguess is a web (war) example containing only non-transactional managed beans. This example can be run on a wide range of servers, including JBoss AS, GlassFish, Apache Tomcat, Jetty, Google App Engine, and any compliant Java EE 6 container. Translator is an enterprise (ear) example that contains session beans. This example must be run on JBoss AS 6.0, GlassFish 3.0 or any compliant Java EE 6 container.

Both examples use JSF 2.0 as the web framework and, as such, can be found in the `examples/jsf` directory of the Weld distribution.

6.1. Prerequisites

To run the examples with the provided build scripts, you'll need the following:

- the latest release of Weld, which contains the examples
- Ant 1.7.0, to build and deploy the examples
- a supported runtime environment (minimum versions shown)
 - JBoss AS 6.0.0,
 - GlassFish 3.0,
 - Apache Tomcat 6.0.x (war example only), or
 - Jetty 6.1.x (war example only)
- (optionally) Maven 2.x, to run the examples in an embedded servlet container



Note

You'll need a full install of Ant 1.7.0. Some linux distributions only supply a partial installation of Ant which cause the build to fail. If you encounter problems, verify that `ant-nodeps.jar` is on the classpath.

In the next few sections, you'll be using the Ant command (`ant`) to invoke the Ant build script in each example to compile, assemble and deploy the example to JBoss AS and, for the war example, Apache Tomcat. You can also deploy the generated artifact (war or ear) to any other container that supports Java EE 6, such as GlassFish 3.

If you have Maven installed, you can use the Maven command (`mvn`) to compile and assemble the standalone artifact (war or ear) and, for the war example, run it in an embedded container.

The sections below cover the steps for deploying with both Ant and Maven in detail. Let's start with JBoss AS.

6.2. Deploying to JBoss AS

To deploy the examples to JBoss AS, you'll need [JBoss AS 6.0.0](http://jboss.org/jbossas/) [http://jboss.org/jbossas/] or above. If a release of the JBoss AS 6.0 line isn't yet available, you can download a [nightly snapshot](http://hudson.jboss.org/hudson/view/JBoss%20AS/job/JBoss-AS-6.0.x/) [http://hudson.jboss.org/hudson/view/JBoss%20AS/job/JBoss-AS-6.0.x/]. The reason JBoss AS 6.0.0 or above is required is because it's the first release

that has both CDI and Bean Validation support built-in, making it close enough to Java EE 6 to run the examples. The good news is that there are no additional modifications you have to make to the server. It's ready to go!

After you have downloaded JBoss AS, extract it. (We recommended renaming the folder to include the `as` qualifier so it's clear that it's the application server). You can move the extracted folder anywhere you like. Wherever it lays to rest, that's what we'll call the JBoss AS installation directory, or `JBOSS_HOME`.

```
$> unzip jboss-6.0.*.zip
$> mv jboss-6.0.* jboss-as-6.0
```

In order for the build scripts to know where to deploy the example, you have to tell them where to find your JBoss AS installation (i.e., `JBOSS_HOME`). Create a new file named `local.build.properties` in the `examples` directory of the Weld distribution and assign the path of your JBoss AS installation to the property key `jboss.home`, as follows:

```
jboss.home=/path/to/jboss-as-6.0
```

You're now ready to deploy your first example!

Switch to the `examples/jsf/numberguess` directory and execute the Ant `deploy` target:

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

If you haven't already, start JBoss AS. You can either start JBoss AS from a Linux shell:

```
$> cd /path/to/jboss-as-6.0
$> ./bin/run.sh
```

a Windows command window:

```
$> cd c:\path\to\jboss-as-6.0\bin
$> run
```

or you can start the server using an IDE, like Eclipse.



Note

If you are using Eclipse, you should seriously consider installing the [JBoss Tools](http://www.jboss.org/tools) [http://www.jboss.org/tools] add-ons, which include a wide variety of tooling for JSR-299 and Java EE development, as well as an enhanced JBoss AS server view.

Wait a few seconds for the application to deploy (or the application server to start) and see if you can determine the most efficient approach to pinpoint the random number at the local URL <http://localhost:8080/weld-numberguess>.



Note

The Ant build script includes additional targets for JBoss AS to deploy and undeploy the archive in either exploded or packaged format and to tidy things up.

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

The second starter example, `weld-translator`, will translate your text into Latin. (Well, not really, but the stub is there for you to implement, at least. Good luck!) To try it out, switch to the translator example directory and execute the deploy target:

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



Note

The translator uses session beans, which are packaged in an EJB module within an ear. Java EE 6 will allow session beans to be deployed in war modules, but that's a topic for a later chapter.

Again, wait a few seconds for the application to deploy (if you're really bored, read the log messages), and visit <http://localhost:8080/weld-translator> to begin pseudo-translating.

6.3. Deploying to GlassFish

Deploying to GlassFish should be easy and familiar, right? After all, it's the Java EE 6 reference implementation and Weld is the JSR-299 reference implementation, meaning Weld gets bundled with GlassFish. So yes, it's all quite easy and familiar.

To deploy the examples to GlassFish, you'll need the final [GlassFish V3](https://glassfish.dev.java.net/downloads/v3-final.html) [https://glassfish.dev.java.net/downloads/v3-final.html] release. Select the release that ends in either `-unix.sh` or `-windows.exe` depending on your platform. After the download is complete, execute the installer. On Linux/Unix, you'll need to first make the script executable.

```
$> chmod 755 glassfish-v3-unix.sh
$> ./glassfish-v3-unix.sh
```

On Windows you can just click on the executable. Follow the instructions in the installer. It will create a single domain named `domain1`. You'll use that domain to deploy the example. We recommend that you choose `7070` as the main HTTP port to avoid conflicts with a running instance of JBoss AS (or Apache Tomcat).

If you've deployed either of the starter examples, `weld-numberguess` or `weld-translator`, to JBoss AS, then you already have the deployable artifact you need. If not, switch to either of the two directories and build it.

```
$> cd examples/jsf/numberguess (or examples/jsf/translator)
$> ant package
```

The deployable archive for the `weld-numberguess`, named `weld-numberguess.war`, ends up in the example's `target` directory. The archive for the `weld-translator` example, named `weld-translator.ear`, ends up in the example's `ear/target` directory. All you need to do now is deploy them to GlassFish.

One way to deploy applications to GlassFish is by using the [GlassFish Admin Console](http://localhost:4848) [http://localhost:4848]. To get the Admin Console running, you need to start a GlassFish domain, in our case `domain1`. Switch to the `bin` folder in the directory where you installed GlassFish and execute the following command:

```
$> asadmin start-domain domain1
```

After a few seconds you can visit the Admin Console in the browser at the URL <http://localhost:4848>. In the tree on the left-hand side of the page, click on "Applications", then click on the "Deploy..." button under the heading "Applications" and select the deployable artifact for either of the two examples. The deployer should recognize that you have selected a Java EE artifact and allow you to start it. You can see the examples running at either <http://localhost:7070/weld-numberguess> or <http://localhost:7070/weld-translator>, depending on which example you deployed.

Alternatively, you can deploy the application to GlassFish using the `asadmin` command:

```
$> asadmin deploy target/weld-numberguess.war
```

The reason the same artifact can be deployed to both JBoss AS and GlassFish, without any modifications, is because all of the features being used are part of the standard platform. And what a capable platform it has become!

6.4. Deploying to Apache Tomcat

Servlet containers are not required to support Java EE services like CDI. However, you can use CDI in a servlet container like Tomcat by embedding a standalone CDI implementation such as Weld.

Weld comes with a servlet listener which bootstraps the CDI environment, registers the `BeanManager` in JNDI and provides injection into servlets. Basically, it emulates some of the work done by the Java EE container. (But you don't get enterprise features such as session beans and container-managed transactions.)

Let's give the Weld servlet extension a spin on Apache Tomcat. First, you'll need to download Tomcat 6.0.18 or later from tomcat.apache.org [http://tomcat.apache.org/download-60.cgi] and extract it.

```
$> unzip apache-tomcat-6.0.18.zip
```

You have two choices for how you can deploy the application to Tomcat. You can deploy it by pushing the artifact to the hot deploy directory using Ant or you can deploy to the server across HTTP using a Maven plugin. The Ant

approach doesn't require that you have Maven installed, so we'll start there. If you want to use Maven, you can just skip ahead.

6.4.1. Deploying with Ant

In order for Ant to push the artifact to the Tomcat hot deploy directory, it needs to know where the Tomcat installation is located. Again, we need to set a property in the `local.build.properties` file in the `examples` directory of the Weld distribution. If you haven't yet created this file, do so now. Then assign the path of your Tomcat installation to the property key `tomcat.home`.

```
tomcat.home=/path/to/apache-tomcat-6
```

Now you're ready to deploy the `numberguess` example to Tomcat!

Change to the `examples/jsf/numberguess` directory again and run the Ant `deploy` target for Tomcat:

```
$> cd examples/jsf/numberguess
$> ant tomcat.deploy
```



Note

The Ant build script includes additional targets for Tomcat to deploy and undeploy the archive in either exploded or packaged format. They are the same target names used for JBoss AS, prefixed with "tomcat.".

- `ant tomcat.restart` - deploy the example in exploded format to Tomcat
- `ant tomcat.explode` - update an exploded example, without restarting the deployment
- `ant tomcat.deploy` - deploy the example in compressed jar format to Tomcat
- `ant tomcat.undeploy` - remove the example from Tomcat

If you haven't already, start Tomcat. You can either start Tomcat from a Linux shell:

```
$> cd /path/to/apache-tomcat-6
$> ./bin/start.sh
```

a Windows command window:

```
$> cd c:\path\to\apache-tomcat-6\bin
$> start
```

or you can start the server using an IDE, like Eclipse.

Wait a few seconds for the application to deploy (or the application server to start) and see if you can figure out the most efficient approach to pinpoint the random number at the local URL <http://localhost:8080/weld-numberguess!>

6.4.2. Deploying with Maven

You can also deploy the application to Tomcat using Maven. This section is a bit more advanced, so skip it unless you're itching to use Maven natively. Of course, you'll first need to make sure that you have Maven installed on your path, similar to how you setup Ant.

The Maven plugin communicates with Tomcat over HTTP, so it doesn't care where you have installed Tomcat. However, the plugin configuration assumes you are running Tomcat in its default configuration, with a hostname of localhost and port 8080. The `readme.txt` file in the example directory has information about how to modify the Maven settings to accommodate a different setup.

To allow Maven to communicate with Tomcat over HTTP, edit the `conf/tomcat-users.xml` file in your Tomcat installation and add the following line:

```
<user username="admin" password="" roles="manager" />
```

Restart Tomcat. You can now deploy the application to Tomcat with Maven using this command:

```
$> mvn compile war:exploded tomcat:exploded -Ptomcat
```

Once the application is deployed, you can redeploy it using this command:

```
$> mvn tomcat:redeploy -Ptomcat
```

The `-Ptomcat` argument activates the `tomcat` profile defined in the Maven POM (`pom.xml`). Among other things, this profile activates the Tomcat plugin.

Rather than shipping the container off to a standalone Tomcat installation, you can also execute the application in an embedded Tomcat 6 container:

```
$> mvn war:inplace tomcat:run -Ptomcat
```

The advantage of using the embedded server is that changes to assets in `src/main/webapp` take effect immediately. If a change to a webapp configuration file is made, the application may automatically redeploy (depending on the plugin configuration). If you make a change to a classpath resource, you need to execute a build:

```
$> mvn compile war:inplace -Ptomcat
```

There are several other Maven goals that you can use if you are hacking on the example, which are documented in the example's `readme.txt` file.

6.5. Deploying to Jetty

Support for Jetty in the examples is a more recent addition. Since Jetty is traditionally used with Maven, there are no Ant targets. You must invoke the Maven build directly to deploy the examples to Jetty out of the box. Also, only the `weld-numberguess` example is configured for Jetty support at the time of writing.

If you've read through the entire Tomcat section, then you're all ready to go. The Maven build parallels the embedded Tomcat deployment. If not, don't worry. We'll still go over everything that you need to know again in this section.

The Maven POM (`pom.xml`) includes a profile named `jetty` that activates the Maven Jetty plugin, which you can use to start Jetty in embedded mode and deploy the application in place. You don't need anything else installed except to have the Maven command (`mvn`) on your path. The rest will be downloaded from the internet when the build is run.

To run the `weld-numberguess` example on Jetty, switch to the example directory and execute the `inplace` goal of the Maven war plugin followed by the `run` goal of the Maven Jetty plugin with the `jetty` profile enabled, as follows:

```
$> cd examples/jsf/numberguess
$> mvn war:inplace jetty:run -Pjetty
```

The log output of Jetty will be shown in the console. Once Jetty reports that the application has deployed, you can access it at the following local URL: <http://localhost:9090/weld-numberguess>. The port is defined in the Maven Jetty plugin configuration within the `jetty` profile.

Any changes to assets in `src/main/webapp` take effect immediately. If a change to a webapp configuration file is made, the application may automatically redeploy. The redeploy behavior can be fined-tuned in the plugin configuration. If you make a change to a classpath resource, you need to execute a build and the `inplace` goal of the Maven war plugin, again with the `jetty` profile enabled.

```
$> mvn compile war:inplace -Pjetty
```

The `war:inplace` goal copies the compiled classes and jars inside `src/main/webapp`, under `WEB-INF/classes` and `WEB-INF/lib`, respectively, mixing source and compiled files. However, the build does work around these temporary files by excluding them from the packaged war and cleaning them during the Maven clean phase.

You have two options if you want to run the example on Jetty from the IDE. You can either install the `m2eclipse` plugin and run the goals as described above. Your other option is to start the Jetty container from a Java application.

First, initialize the Eclipse project:

```
$> mvn clean eclipse:clean eclipse:eclipse -Pjetty-ide
```

Next, assemble all the necessary resources under `src/main/webapp`:

```
$> mvn war:inplace -Pjetty-ide
```

Now, you are ready to run the server in Eclipse. Import the project into your Eclipse workspace using "Import Existing Project into Workspace". Then, find the start class in `src/jetty/java` and run its main method as a Java Application. Jetty will launch. You can view the application at the following local URL: <http://localhost:8080>. Pay particular attention to the port in the URL and the lack of a trailing context path.

Now that you have gotten the starter applications deployed on the server of your choice, you probably want to know a little bit about how they actually work.

Diving into the Weld examples

It's time to pull the covers back and dive into the internals of Weld example applications. Let's start with the simpler of the two examples, `weld-numberguess`.

7.1. The numberguess example in depth

In the `numberguess` application you get 10 attempts to guess a number between 1 and 100. After each attempt, you're told whether your guess was too high or too low.

The `numberguess` example is comprised of a number of beans, configuration files and Facelets (JSF) views, packaged as a war module. Let's start by examining the configuration files.

All the configuration files for this example are located in `WEB-INF/`, which can be found in the `src/main/webapp` directory of the example. First, we have the JSF 2.0 version of `faces-config.xml`. A standardized version of Facelets is the default view handler in JSF 2.0, so there's really nothing that we have to configure. Thus, the configuration consists of only the root element.

```
<faces-config version="2.0"
  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_2_0.xsd">
</faces-config>
```

There's also an empty `beans.xml` file, which tells the container to look for beans in this application and to activate the CDI services.

Finally, there's the familiar `web.xml`:

```
<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>weld-jsf-numberguess-war</display-name>
  <description>Weld JSF numberguess example (war)</description>

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

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

1

2

```

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

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

</web-app>

```

- ❶ Enable and initialize the JSF servlet
- ❷ Configure requests for URLs ending in `.jsf` to be handled by JSF
- ❸ Tell JSF that we will be giving our JSF views (Facelets templates) an extension of `.xhtml`
- ❹ Configure a session timeout of 10 minutes



Note

This demo uses JSF 2 as the view framework, but you can use Weld with any servlet-based web framework, such as JSF 1.2 or Wicket.

Let's take a look at the main JSF view, `src/main/webapp/home.xhtml`.

```

<!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">

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

        <div style="color: red">
          <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.biggest}.
          You have #{game.remainingGuesses} guesses remaining.
        </div>

        <div>
          Your guess:

```

```

        <h:inputText id="inputGuess" value="#{game.guess}"
                    size="3" required="true" disabled="#{game.number eq game.guess}"
                    validator="#{game.validateNumberRange}" />
        <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>

```

- ① Facelets is the built-in templating language for JSF. Here we are wrapping our page in a template which defines the layout.
- ② There are a number of messages which can be sent to the user, "Higher!", "Lower!" and "Correct!"
- ③ As the user guesses, the range of numbers they can guess gets smaller - this sentence changes to make sure they know the number range of a valid guess.
- ④ This input field is bound to a bean property using a value expression.
- ⑤ A validator binding 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 bounds number.
- ⑥ 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 bean.

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

```

@Qualifier
@Target( { TYPE, METHOD, PARAMETER, FIELD })
@Retention(RUNTIME)
public @interface Random {}

```

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

```

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

```

The application-scoped `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 implements Serializable {

    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;
}
}

```

The Generator is application scoped, so we don't get a different random each time.



Note

The package declaration and imports have been excluded from these listings. The complete listing is available in the example source code.

The final bean in the application is the session-scoped Game class. This is the primary entry point of the application. It's responsible for setting up or resetting the game, capturing and validating the user's guess and providing feedback to the user with a FacesMessage. We've used the post-construct lifecycle method to initialize the game by retrieving a random number from the @Random Instance<Integer> bean.

You'll notice that we've also added the @Named annotation to this class. This annotation is only required when you want to make the bean accessible to a JSF view via EL (i.e., #{game}).

```

@Named
@SessionScoped
public class Game implements Serializable {

    private int number;
    private int guess;
    private int smallest;
    private int biggest;
    private int remainingGuesses;

    @Inject @MaxNumber private int maxNumber;
    @Inject @Random Instance<Integer> randomNumber;

    public Game() {}

    public void check() {
        if (guess > number) {
            biggest = guess - 1;
        }
        else if (guess < number) {
            smallest = guess + 1;
        }
    }
}

```

```

        else if (guess == number) {
            FacesContext.getCurrentInstance().addMessage(null, new FacesMessage("Correct!"));
        }
        remainingGuesses--;
    }

    @PostConstruct
    public void reset() {
        this.smallest = 0;
        this.guess = 0;
        this.remainingGuesses = 10;
        this.biggest = maxNumber;
        this.number = randomNumber.get();
    }

    public void validateNumberRange(FacesContext context, UIComponent toValidate, Object value) {
        if (remainingGuesses <= 0) {
            FacesMessage message = new FacesMessage("No guesses left!");
            context.addMessage(toValidate.getClientId(context), message);
            ((UIInput) toValidate).setValid(false);
            return;
        }
        int input = (Integer) value;

        if (input < smallest || input > biggest) {
            ((UIInput) toValidate).setValid(false);

            FacesMessage message = new FacesMessage("Invalid guess");
            context.addMessage(toValidate.getClientId(context), message);
        }
    }

    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;
    }
}

```

7.1.1. The numberguess example in Apache Tomcat or Jetty

A couple of modifications must be made to the numberguess artifact in order to deploy it to Tomcat or Jetty. First, Weld must be deployed as a Web Application library under `WEB-INF/lib` since the servlet container does not provide the CDI services. For your convenience we provide a single jar suitable for running Weld in any servlet container (including Jetty), `weld-servlet.jar`.



Note

You must also include the jars for JSF, EL, and the common annotations (`jsr250-api.jar`), all of which are provided by the Java EE platform (a Java EE application server). Are you starting to appreciate why a Java EE platform is worth using?

Second, we need to explicitly specify the servlet listener in `web.xml`, again because the container isn't doing this stuff for you. The servlet listener boots Weld and controls its interaction with requests.

```
<listener>
  <listener-class>org.jboss.weld.environment.servlet.Listener</listener-class>
</listener>
```

When Weld boots, it places the `javax.enterprise.inject.spi.BeanManager`, the portable SPI for obtaining bean instances, in the `ServletContext` under a variable name equal to the fully-qualified interface name. You generally don't need to access this interface, but Weld makes use of it.

7.2. The numberguess example for Java SE with Swing

This example shows how to use the Weld SE extension to in a Java SE based Swing application with no EJB or servlet dependencies. This example can be found in the `examples/se/numberguess` folder of the Weld distribution.

7.2.1. Creating the Eclipse project

To use the Weld SE numberguess example in Eclipse, you can open the example natively using the [m2eclipse plugin](http://m2eclipse.sonatype.org/) [http://m2eclipse.sonatype.org/].

If you have m2eclipse installed, you can open any Maven project directly. From within Eclipse, select *File -> Import... -> Maven Projects*. Then, browse to the location of the Weld SE numberguess example. You should see that Eclipse recognizes the existence of a Maven project.

This will create a project in your workspace called `weld-se-numberguess`.

If you are not using the m2eclipse plugin, you have to follow different steps to import the project. First, switch into the Weld SE numberguess example, then execute the Maven Eclipse plugin with the jetty profile activated, as follows:

It's time to get the example running!

7.2.2. Running the example from Eclipse

Disable m2eclipse's *Workspace Resolution*, to make sure that Eclipse can find `StartMain`. Right click on the project, and choose *Properties -> Maven*, and uncheck *Resolve dependencies from Workspace projects*:



Right click on the project, and choose *Run As -> Run As Application*:



Locate the StartMain class:



The application should now launch!



7.2.3. Running the example from the command line

- Ensure that Maven 3 is installed and in your PATH
- Ensure that the `JAVA_HOME` environment variable is pointing to your JDK installation
- Open a command line or terminal window in the `examples/se/numberguess` directory
- Execute the following command

```
mvn -Drun
```

7.2.4. Understanding the code

Let's have a look at the significant code and configuration files that make up this example.

As usual, there is an empty `beans.xml` file in the root package (`src/main/resources/beans.xml`), which marks this application as a CDI application.

The game's main logic is located in `Game.java`. Here is the code for that class, highlighting the ways in which this differs from the web application version:

```
@ApplicationScoped
public class Game
{
    public static final int MAX_NUM_GUESSES = 10;

    private Integer number;
    private int guess = 0;
    private int smallest = 0;

    @Inject
    @MaxNumber
    private int maxNumber;

    private int biggest;
    private int remainingGuesses = MAX_NUM_GUESSES;
    private boolean validNumberRange = true;

    @Inject
    Generator rndGenerator;

    public Game()
    {
    }

    ...

    public boolean isValidNumberRange()
    {
        return validNumberRange;
    }

    public boolean isGameWon()
    {
        return guess == number;
    }

    public boolean isGameLost()
    {
        return guess != number && remainingGuesses <= 0;
    }

    public boolean check()
    {
        boolean result = false;

        if (checkNewNumberRangeIsValid())
        {
            if (guess > number)
            {
                biggest = guess - 1;
            }

            if (guess < number)
            {
                smallest = guess + 1;
            }
        }
    }
}
```

```

    }

    if (guess == number)
    {
        result = true;
    }

    remainingGuesses--;
}

return result;
}

private boolean checkNewNumberRangeIsValid()
{
    return validNumberRange = ((guess >= smallest) && (guess <= biggest));
}

@PostConstruct
public void reset()
{
    this.smallest = 0;
    this.guess = 0;
    this.remainingGuesses = 10;
    this.biggest = maxNumber;
    this.number = rndGenerator.next();
}
}

```

- ① The bean is application scoped rather than session scoped, since an instance of a Swing application typically represents a single 'session'.
- ② Notice that the bean is not named, since it doesn't need to be accessed via EL.
- ③ In Java SE there is no JSF `FacesContext` to which messages can be added. Instead the `Game` class provides additional information about the state of the current game including:

- If the game has been won or lost
- If the most recent guess was invalid

This allows the Swing UI to query the state of the game, which it does indirectly via a class called `MessageGenerator`, in order to determine the appropriate messages to display to the user during the game.

- ④ Since there is no dedicated validation phase, validation of user input is performed during the `check()` method.
- ⑤ The `reset()` method makes a call to the injected `rndGenerator` in order to get the random number at the start of each game. Note that it can't use `Instance.get()` like the JSF example does because there will not be any active contexts like there are during a JSF request.

The `MessageGenerator` class depends on the current instance of `Game` and queries its state in order to determine the appropriate messages to provide as the prompt for the user's next guess and the response to the previous guess. The code for `MessageGenerator` is as follows:

```

public class MessageGenerator
{
    @Inject

```

```

private Game game;

public String getChallengeMessage() 2
{
    StringBuilder challengeMsg = new StringBuilder("I'm thinking of a number between ");
    challengeMsg.append(game.getSmallest());
    challengeMsg.append(" and ");
    challengeMsg.append(game.getBiggest());
    challengeMsg.append(". Can you guess what it is?");

    return challengeMsg.toString();
}

public String getResultMessage() 3
{
    if (game.isGameWon())
    {
        return "You guessed it! The number was " + game.getNumber();
    }
    else if (game.isGameLost())
    {
        return "You are fail! The number was " + game.getNumber();
    }
    else if (!game.isValidNumberRange())
    {
        return "Invalid number range!";
    }
    else if (game.getRemainingGuesses() == Game.MAX_NUM_GUESSES)
    {
        return "What is your first guess?";
    }
    else
    {
        String direction = null;

        if (game.getGuess() < game.getNumber())
        {
            direction = "Higher";
        }
        else
        {
            direction = "Lower";
        }

        return direction + "! You have " + game.getRemainingGuesses() + " guesses left.";
    }
}
}

```

- 1 The instance of Game for the application is injected here.
- 2 The Game's state is interrogated to determine the appropriate challenge message ...
- 3 ... and again to determine whether to congratulate, console or encourage the user to continue.

Finally we come to the NumberGuessFrame class which provides the Swing front end to our guessing game.

```

public class NumberGuessFrame extends javax.swing.JFrame
{
    @Inject
    private Game game;

    @Inject
    private MessageGenerator msgGenerator;

    public void start(@Observes ContainerInitialized event)
    {
        java.awt.EventQueue.invokeLater(new Runnable()
        {
            public void run()
            {
                initComponents();
                setVisible(true);
            }
        });
    }

    private void initComponents()
    {
        buttonPanel = new javax.swing.JPanel();
        mainMsgPanel = new javax.swing.JPanel();
        mainLabel = new javax.swing.JLabel();
        messageLabel = new javax.swing.JLabel();
        guessText = new javax.swing.JTextField();
        ...
        mainLabel.setText(msgGenerator.getChallengeMessage());
        mainMsgPanel.add(mainLabel);

        messageLabel.setText(msgGenerator.getResultMessage());
        mainMsgPanel.add(messageLabel);
        ...
    }

    private void guessButtonActionPerformed( java.awt.event.ActionEvent evt )
    {
        int guess = Integer.parseInt(guessText.getText());
        game.setGuess( guess );
        game.check();
        refreshUI();
    }

    private void replayBtnActionPerformed(java.awt.event.ActionEvent evt)
    {
        game.reset();
        refreshUI();
    }

    private void refreshUI() {
        mainLabel.setText( msgGenerator.getChallengeMessage() );
        messageLabel.setText( msgGenerator.getResultMessage() );
        guessText.setText( " " );
    }
}

```

```

        guessesLeftBar.setValue( game.getRemainingGuesses() );
        guessText.requestFocus();
    }

    // swing components
    private javax.swing.JPanel borderPanel;
    ...
    private javax.swing.JButton replayBtn;
}

```

- ① The injected instance of the game (logic and state).
- ② The injected message generator for UI messages.
- ③ This application is started in the prescribed Weld SE way, by observing the `ContainerInitialized` event.
- ④ This method initializes all of the Swing components. Note the use of the `msgGenerator` here.
- ⑤ `guessButtonActionPerformed` is called when the 'Guess' button is clicked, and it does the following:
 - Gets the guess entered by the user and sets it as the current guess in the `Game`
 - Calls `game.check()` to validate and perform one 'turn' of the game
 - Calls `refreshUI`. If there were validation errors with the input, this will have been captured during `game.check()` and as such will be reflected in the messages returned by `MessageGenerator` and subsequently presented to the user. If there are no validation errors then the user will be told to guess again (higher or lower) or that the game has ended either in a win (correct guess) or a loss (ran out of guesses).
- ⑥ `replayBtnActionPerformed` simply calls `game.reset()` to start a new game and refreshes the messages in the UI.

7.3. The translator example in depth

The translator example will take any sentences you enter, and translate them to Latin. (Well, not really, but the stub is there for you to implement, at least. Good luck!)

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



Note

Java EE 6, which bundles EJB 3.1, allows you to package EJBs in a war, which will make this structure much simpler! Still, there are other advantages of using an ear.

First, let's take a look at the ear aggregator, which is located in the example's `ear` directory. Maven automatically generates the `application.xml` for us from this plugin configuration:

```

<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-ear-plugin</artifactId>
  <configuration>
    <modules>
      <webModule>
        <groupId>org.jboss.weld.examples.jsf.translator</groupId>

```

```

        <artifactId>weld-jsf-translator-war</artifactId>
        <contextRoot>/weld-translator</contextRoot>
    </webModule>
</modules>
</configuration>
</plugin>

```

This configuration overrides the web context path, resulting in this application URL: <http://localhost:8080/weld-translator>.



Note

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

```

<application version="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/application_5.xsd">

  <display-name>weld-jsf-translator-ear</display-name>
  <description>The Weld JSF translator example (ear)</description>

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

```

Next, let's look at the war, which is located in the example's war directory. Just as in the numberguess example, we have a faces-config.xml for JSF 2.0 and a web.xml (to activate JSF) under WEB-INF, both sourced from src/main/webapp/WEB-INF.

More interesting is the JSF view 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="translator">

  <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 left-hand textarea, and hit the translate button to see the result to the right.

Finally, let's look at the EJB module, which is located in the example's `ejb` directory. In `src/main/resources/META-INF` there is just an empty `beans.xml`, used to mark the archive as containing beans.

We've saved the most interesting bit to last, the code! The project has two simple beans, `SentenceParser` and `TextTranslator` and two session beans, `TranslatorControllerBean` and `SentenceTranslator`. You should be getting quite familiar with what a 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 injection:

```

public class TextTranslator implements Serializable {

    private SentenceParser sentenceParser;

    @EJB private Translator translator;

    @Inject public TextTranslator(SentenceParser sentenceParser) {
        this.sentenceParser = sentenceParser;
    }

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

}

```

`TextTranslator` uses the simple bean (really just a plain Java class!) `SentenceParser` to parse the sentence and then calls on the stateless bean with the local business interface `Translator` to perform the translation. That's where the magic happens. Of course, we couldn't develop a full translator, but it's convincing enough to anyone who doesn't understand Latin!


```
@Stateless
public class SentenceTranslator implements Translator {

    public String translate(String sentence) {
        return "Lorem ipsum dolor sit amet";
    }

}
```

Finally, there is UI orientated controller. This is a request scoped, named, stateful session bean, which injects the translator. It collects the text from the user and dispatches it to the translator. The bean also has getters and setters for all the fields on the page.

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

    @Inject private TextTranslator translator;

    private String inputText;

    private String translatedText;

    public void translate() {
        translatedText = translator.translate(inputText);
    }

    public String getText() {
        return inputText;
    }

    public void setText(String text) {
        this.inputText = text;
    }

    public String getTranslatedText() {
        return translatedText;
    }

    @Remove public void remove() {}

}
```

That concludes our short tour of the Weld starter examples. For more information on Weld, please visit <http://www.seamframework.org/Weld>.

Part III. Loose coupling with strong typing

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

- *alternatives* enable deployment time polymorphism,
- *producer methods* enable runtime polymorphism, and
- *contextual lifecycle management* decouples 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 interface, nor is it required to manage the lifecycle of the implementation. 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 (most notably by using XML descriptors). CDI is the first technology, and certainly the first specification in the Java EE platform, that achieves this level of loose coupling in a typesafe way.

CDI 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.

The second major theme of CDI is *strong typing*. The information about the dependencies, interceptors and decorators of a 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 CDI 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.

CDI 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.

CDI *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, qualifiers, etc, into a single reusable package. (Of course, there is also the benefit of tucking some of those annotations away).

We're now ready to meet some more advanced features of CDI. 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.

Producer methods

Producer methods let us overcome certain limitations that arise when a container, instead of the application, is responsible for instantiating objects. They're also the easiest way to integrate objects which are not beans into the CDI environment.

According to the spec:

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

- the objects to be injected are not required to be instances of 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 bean constructor

For example, producer methods let us:

- expose a JPA entity as a bean,
- expose any JDK class as a bean,
- define multiple beans, with different scopes or initialization, for the same implementation class, or
- vary the implementation of a bean type at runtime.

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

```
@SessionScoped
public class Preferences implements Serializable {
    private PaymentStrategyType paymentStrategy;
    ...
    @Produces @Preferred
    public PaymentStrategy getPaymentStrategy() {
        switch (paymentStrategy) {
            case CREDIT_CARD: return new CreditCardPaymentStrategy();
            case CHECK: return new CheckPaymentStrategy();
            case PAYPAL: return new PayPalPaymentStrategy();
            default: return null;
        }
    }
}
```

Consider an injection point:

```
@Inject @Preferred PaymentStrategy paymentStrategy;
```

This injection point has the same type and qualifier annotations as the producer method, so it resolves to the producer method using the usual CDI injection rules. The producer method will be called by the container to obtain an instance to service this injection point.

8.1. Scope of a producer method

The scope of the producer method defaults to `@Dependent`, and so it will be called *every time* the container 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.



Note

A producer method does *not* inherit the scope of the bean that declares the method. There are two different beans here: the producer method, and the bean which declares it. The scope of the producer method determines how often the method will be called, and the lifecycle of the objects returned by the method. The scope of the bean that declares the producer method determines the lifecycle of the object upon which the producer method is invoked.

8.2. 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 bean instances:

```
@Produces @Preferred @SessionScoped
public PaymentStrategy getPaymentStrategy(CreditCardPaymentStrategy ccps,
                                          CheckPaymentStrategy 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 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 container 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 container, so please take extra care when returning 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 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` qualifier annotation.

8.3. Use of @New with producer methods

Consider the following producer method:

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

8.4. Disposer methods

Some producer methods return objects that require explicit destruction. For example, somebody needs to close this JDBC connection:

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

Destruction can be performed by a matching *disposer method*, defined by the same class as the producer method:

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

The disposer method must have at least one parameter, annotated `@Disposes`, with the same type and qualifiers as the producer method. The disposer method is called automatically when the context ends (in this case, at the end of the request), and this parameter receives the object produced by the producer method. If the disposer method has additional method parameters, the container will look for a bean that satisfies the type and qualifiers of each parameter and pass it to the method automatically.

Interceptors

Interceptor functionality is defined in the Java Interceptors specification. CDI enhances this functionality with a more sophisticated, semantic, annotation-based approach to binding interceptors to beans.

The Interceptors specification defines two kinds of interception points:

- business method interception, and
- lifecycle callback interception.

In addition, the EJB specification defines timeout method interception.

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

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

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.

A *timeout method interceptor* applies to invocations of EJB timeout methods by the container:

```
public class TimeoutInterceptor {
    @AroundTimeout
    public Object manageTransaction(InvocationContext ctx) throws Exception { ... }
}
```

9.1. Interceptor bindings

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

```
@InterceptorBinding
@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() { ... }
}
```

9.2. 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 interceptor, and annotate it `@Interceptor` and `@Transactional`.

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

Interceptors can take advantage of dependency injection:

```
@Transactional @Interceptor
public class TransactionInterceptor {

    @Resource UserTransaction transaction;

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

}
```

Multiple interceptors may use the same interceptor binding type.

9.3. Enabling interceptors

By default, all interceptors are disabled. We need to *enable* our interceptor in the `beans.xml` descriptor of a bean archive. This activation only applies to the beans in that archive.

```
<beans
    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/beans_1_0.xsd">
    <interceptors>
        <class>org.mycompany.myapp.TransactionInterceptor</class>
    </interceptors>
</beans>

```

Whoah! Why the angle bracket stew?

Well, having the XML declaration is actually a *good thing*. It 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 transaction interceptor.

```

<beans
    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/beans_1_0.xsd">
    <interceptors>
        <class>org.mycompany.myapp.SecurityInterceptor</class>
        <class>org.mycompany.myapp.TransactionInterceptor</class>
    </interceptors>
</beans>

```

Or we could turn them both off in our test environment by simply not mentioning them in `beans.xml`! Ah, so simple.

9.4. Interceptor bindings with members

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

```

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

```

CDI 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) throws Exception { ... }
}

```

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 container to ignore the value of `requiresNew` when binding interceptors? Perhaps this information is only useful for the interceptor implementation. We can use the `@Nonbinding` annotation:

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

9.5. Multiple interceptor binding annotations

Usually we use combinations of interceptor bindings types to bind multiple interceptors to a bean. For example, the following declaration would be used to bind `TransactionInterceptor` and `SecurityInterceptor` to the same 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() { ... }
}
```

9.6. 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, CDI works around this missing feature of Java. We may annotate one interceptor binding type with other interceptor binding types (termed a *meta-annotation*). The interceptor bindings are transitive — any bean with the first interceptor binding inherits the interceptor bindings declared as meta-annotations.

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

Now, any bean annotated `@Action` will be bound to both `TransactionInterceptor` and `SecurityInterceptor`. (And even `TransactionalSecureInterceptor`, if it exists.)

9.7. Use of `@Interceptors`

The `@Interceptors` annotation defined by the interceptor specification (and used by the managed bean and EJB specifications) is still supported in CDI.

```
@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 CDI-style interceptor bindings.

Decorators

Interceptors are a powerful way to capture and separate concerns which are *orthogonal* to the application (and 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, security and call logging. 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. Since decorators directly implement operations with business semantics, it makes them the 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. Interceptors and decorators, though similar in many ways, are complementary. Let's look at some cases where decorators fit the bill.

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 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 bean (possibly even an abstract class) that implements the type it decorates and is annotated `@Decorator`.

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

The decorator implements the methods of the decorated type that it wants to intercept.

```
@Decorator
public abstract class LargeTransactionDecorator
    implements Account {
    @Inject @Delegate @Any Account account;

    @PersistenceContext EntityManager em;

    public void withdraw(BigDecimal amount) {
        ...
    }
}
```

```
public void deposit(BigDecimal amount);  
    ...  
}  
}
```

Unlike other beans, a decorator may be an abstract class. Therefore, 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.

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

10.1. Delegate object

Decorators have a special injection point, called the *delegate injection point*, with the same type as the beans they decorate, and the annotation `@Delegate`. There must be exactly one delegate injection point, which can be a constructor parameter, initializer method parameter or injected field.

```
@Decorator  
public abstract class LargeTransactionDecorator  
    implements Account {  
    @Inject @Delegate @Any Account account;  
    ...  
}
```

A decorator is bound to any bean which:

- has the type of the delegate injection point as a bean type, and
- has all qualifiers that are declared at the delegate injection point.

This delegate injection point specifies that the decorator is bound to all beans that implement `Account`:

```
@Inject @Delegate @Any Account account;
```

A delegate injection point may specify any number of qualifier annotations. The decorator will only be bound to beans with the same qualifiers.

```
@Inject @Delegate @Foreign Account account;
```

The decorator may invoke the delegate object, which has much the same effect as calling `InvocationContext.proceed()` from an interceptor. The main difference is that the decorator can invoke *any* business method on the delegate object.

```
@Decorator  
public abstract class LargeTransactionDecorator  
    implements Account {  
    @Inject @Delegate @Any 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) );
}
}
}

```

10.2. Enabling decorators

By default, all decorators are disabled. We need to *enable* our decorator in the `beans.xml` descriptor of a bean archive. This activation only applies to the beans in that archive.

```

<beans
  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/beans_1_0.xsd">
  <decorators>
    <class>org.mycompany.myapp.LargeTransactionDecorator</class>
  </decorators>
</beans>

```

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.

Events

Dependency injection enables loose-coupling by allowing the implementation of the injected bean type to vary, either a deployment time or runtime. Events go one step further, allowing beans to interact with no compile time dependency at all. Event *producers* raise events that are delivered to event *observers* by the container.

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.

The CDI event notification facility uses more or less the same typesafe approach that we've already seen with the dependency injection service.

11.1. Event payload

The event object carries state from producer to consumer. The event object is nothing more than an instance of a concrete Java class. (The only restriction is that an event type may not contain type variables). An event may be assigned qualifiers, which allows observers to distinguish it from other events of the same type. The qualifiers function like topic selectors, allowing an observer to narrow the set of events it observes.

An event qualifier is just a normal qualifier, defined using `@Qualifier`. Here's an example:

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

11.2. Event observers

An *observer method* is a method of a 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*, in this case `Document`. The event parameter may also specify qualifiers.

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

An observer method need not specify any event qualifiers—in this case it is interested in *only unqualified* events of a particular type. If it does specify qualifiers, it's only interested in events which have those qualifiers.

The observer method may have additional parameters, which are injection points:

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

11.3. Event producers

Event producers fire events using an instance of the parameterized `Event` interface. An instance of this interface is obtained by injection:

```
@Inject @Any Event<Document> documentEvent;
```

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

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

This particular event will be delivered to every observer method that:

- has an event parameter to which the event object (the `Document`) is assignable, and
- specifies no qualifiers.

The container 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 container stops calling observer methods, and the exception is rethrown by the `fire()` method.

Qualifiers can be applied to an event in one of two ways:

- by annotating the `Event` injection point, or
- by passing qualifiers to the `select()` of `Event`.

Specifying the qualifiers at the injection point is far simpler:

```
@Inject @Updated Event<Document> documentUpdatedEvent;
```

Then, every event fired via this instance of `Event` has the event qualifier `@Updated`. The event is delivered to every observer method that:

- has an event parameter to which the event object is assignable, and
- does not have any event qualifier *except* for the event qualifiers that match those specified at the `Event` injection point.

The downside of annotating the injection point is that we can't specify the qualifier dynamically. CDI lets us obtain a qualifier instance by subclassing the helper class `AnnotationLiteral`. That way, we can pass the qualifier to the `select()` method of `Event`.

```
documentEvent.select(new AnnotationLiteral<Updated>()).fire(document);
```

Events can have multiple event qualifiers, assembled using any combination of annotations at the `Event` injection point and qualifier instances passed to the `select()` method.

11.4. Conditional observer methods

By default, if there is no instance of an observer in the current context, the container will instantiate the observer in order to deliver an event to it. This behavior isn't always desirable. We may want to deliver events only to instances of the observer that already exist in the current contexts.

A conditional observer is specified by adding `receive = IF_EXISTS` to the `@Observes` annotation.

```
public void refreshOnDocumentUpdate(@Observes(receive = IF_EXISTS) @Updated Document d) { ... }
```

A bean with scope `@Dependent` cannot be a conditional observer, since it would never be called!

11.5. Event qualifiers with members

An event qualifier type may have annotation members:

```
@Qualifier
@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 qualifier type members may be specified statically by the event producer, via annotations at the event notifier injection point:

```
@Inject @Role(ADMIN) Event<LoggedIn> loggedInEvent;
```

Alternatively, the value of the event qualifier 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 `select()`:

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

11.6. Multiple event qualifiers

Event qualifier types may be combined, for example:

```
@Inject @Blog Event<Document> blogEvent;  
...  
if (document.isBlog()) blogEvent.select(new AnnotationLiteral<Updated>({})).fire(document);
```

Observers must completely match the final qualified type of the event. Assume the following observers in this example:

```
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) { ... }}
```

The only observer notified will be:

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

However, if there were also an observer:

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

It would also be notified, as `@Any` acts as an alias for any and all qualifiers.

11.7. 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(@Observes(during = AFTER_SUCCESS) CategoryUpdateEvent event) { ... }
```

There are five kinds of transactional observers:

- `IN_PROGRESS` observers are called immediately (default)

`AFTER_SUCCESS` observers are called during the after completion phase of the transaction, but only if the transaction completes successfully

- `AFTER_FAILURE` observers are called during the after completion phase of the transaction, but only if the transaction fails to complete successfully
- `AFTER_COMPLETION` observers are called during the after completion phase of the transaction
- `BEFORE_COMPLETION` observers are called during the before completion phase of the transaction

Transactional observers are very important in a stateful object model 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 bean that creates and deletes `Products` could raise events, for example:

```
@Stateless
public class ProductManager {

    @PersistenceContext EntityManager em;
    @Inject @Any Event<Product> productEvent;

    public void delete(Product product) {
        em.delete(product);
        productEvent.select(new AnnotationLiteral<Deleted>({})).fire(product);
    }
}
```

```
public void persist(Product product) {  
    em.persist(product);  
    productEvent.select(new AnnotationLiteral<Created>()).fire(product);  
}  
...  
}
```

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

```
@ApplicationScoped @Singleton  
public class Catalog {  
    ...  
    void addProduct(@Observes(during = AFTER_SUCCESS) @Created Product product) {  
        products.add(product);  
    }  
  
    void removeProduct(@Observes(during = AFTER_SUCCESS) @Deleted Product product) {  
        products.remove(product);  
    }  
}
```


Stereotypes

The CDI specification defines a stereotype as follows:

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

A stereotype encapsulates any combination of:

- a default scope, and
- a set of interceptor bindings.

A stereotype may also specify that:

- all beans with the stereotype have defaulted bean EL names, or that
- all beans with the stereotype are alternatives.

A bean may declare zero, one or multiple stereotypes. Stereotype annotations may be applied to a bean class or producer method or field.

A stereotype is an annotation, annotated `@Stereotype`, that packages several other annotations. For instance, the following stereotype identifies action classes in some MVC framework:

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

We use the stereotype by applying the annotation to a bean.

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

Of course, we need to apply some other annotations to our stereotype or else it wouldn't be adding much value.

12.1. Default scope for a stereotype

A stereotype may specify a default scope for beans annotated with the stereotype. For example:

```
@RequestScoped
@Stereotype
@Retention(RUNTIME)
@Target(TYPE)
```

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

A particular action may still override this default if necessary:

```
@Dependent @Action
public class DependentScopedLoginAction { ... }
```

Naturally, overriding a single default isn't much use. But remember, stereotypes can define more than just the default scope.

12.2. Interceptor bindings for stereotypes

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

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

This helps us get technical concerns, like transactions and security, even further away from the business code!

12.3. Name defaulting with stereotypes

We can specify that all beans with a certain stereotype have a defaulted EL name when a name is not explicitly defined for that bean. All we need to do is add an empty `@Named` annotation:

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

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

12.4. Alternative stereotypes

A stereotype can indicate that all beans to which it is applied are `@Alternatives`. An *alternative stereotype* lets us classify beans by deployment scenario.

```
@Alternative
@Stereotype
```

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

We can apply an alternative stereotype to a whole set of beans, and activate them all with one line of code in `beans.xml`.

```
@Mock
public class MockLoginAction extends LoginAction { ... }
```

12.5. Stereotype stacking

This may blow your mind a bit, but stereotypes may declare other stereotypes, which we'll call *stereotype stacking*. You may want to do this if you have two distinct stereotypes which are meaningful on their own, but in other situation may be meaningful when combined.

Here's an example that combines the `@Action` and `@Auditable` stereotypes:

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

12.6. Built-in stereotypes

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

CDI defines one further standard stereotype, `@Model`, which is expected to be used frequently in web applications:

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

Instead of using JSF managed beans, just annotate a bean `@Model`, and use it directly in your JSF view!

Specialization, inheritance and alternatives

When you first start developing with CDI, you'll likely be dealing only with a single bean implementation for each bean type. In this case, it's easy to understand how beans get selected for injection. As the complexity of your application grows, multiple occurrences of the same bean type start appearing, either because you have multiple implementations or two beans share a common (Java) inheritance. That's when you have to begin studying the specialization, inheritance and alternative rules to work through unsatisfied or ambiguous dependencies or to avoid certain beans from being called.

The CDI specification recognizes two distinct scenarios in which one bean extends another:

- The second bean *specializes* the first bean in certain deployment scenarios. In these deployments, the second bean completely replaces the first, fulfilling the same role in the system.
- The second bean is simply reusing the Java implementation, and otherwise bears no relation to the first bean. The first bean may not even have been designed for use as a contextual object.

The second case is the default assumed by CDI. It's possible to have two beans in the system with the same part bean type (interface or parent class). As you've learned, you select between the two implementations using qualifiers.

The first case is the exception, and also requires more care. In any given deployment, only one bean can fulfill a given role at a time. That means one bean needs to be enabled and the other disabled. There are two modifiers involved: `@Alternative` and `@Specializes`. We'll start by looking at alternatives and then show the guarantees that specialization adds.

13.1. Using alternative stereotypes

CDI lets you *override* the implementation of a bean type at deployment time using an alternative. For example, the following bean provides a default implementation of the `PaymentProcessor` interface:

```
public class DefaultPaymentProcessor
    implements PaymentProcessor {
    ...
}
```

But in our staging environment, we don't really want to submit payments to the external system, so we override that implementation of `PaymentProcessor` with a different bean:

```
public @Alternative
class StagingPaymentProcessor
    implements PaymentProcessor {
    ...
}
```

or

```

public @Alternative
class StagingPaymentProcessor
    extends DefaultPaymentProcessor {
    ...
}

```

We've already seen how we can enable this alternative by listing its class in the `beans.xml` descriptor.

But suppose we have many alternatives in the staging environment. It would be much more convenient to be able to enable them all at once. So let's make `@Staging` an `@Alternative` stereotype and annotate the staging beans with this stereotype instead. You'll see how this level of indirection pays off. First, we create the stereotype:

```

@Alternative
@Stereotype
@Retention(RUNTIME)
@Target(TYPE)
public @interface Staging {}

```

Then we replace the `@Alternative` annotation on our bean with `@Staging`:

```

@Staging
public class StagingPaymentProcessor
    implements PaymentProcessor {
    ...
}

```

Finally, we activate the `@Staging` stereotype in the `beans.xml` descriptor:

```

<beans
    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/beans_1_0.xsd">
    <alternatives>
        <stereotype>org.mycompany.myapp.Staging</stereotype>
    </alternatives>
</beans>

```

Now, no matter how many staging beans we have, they will all be enabled at once.

13.2. A minor problem with alternatives

When we enable an alternative, does that mean the default implementation is disabled? Well, not exactly. If the default implementation has a qualifier, for instance `@LargeTransaction`, and the alternative does not, you could still inject the default implementation.

```
@Inject @LargeTransaction PaymentProcessor paymentProcessor;
```

So we haven't completely replaced the default implementation in this deployment of the system. The only way one bean can completely override a second bean at all injection points is if it implements all the bean types and declares all the qualifiers of the second bean. However, if the second bean declares a producer method or observer method, then even this is not enough to ensure that the second bean is never called! We need something extra.

CDI provides a special feature, called *specialization*, that helps the developer avoid these traps. Specialization is a way of informing the system of your intent to completely replace and disable an implementation of a bean.

13.3. Using specialization

When the goal is to replace one bean implementation with a second, to help prevent developer error, the first bean may:

- directly extend the bean class of the second bean, or
- directly override the producer method, in the case that the second bean is a producer method, and then

explicitly declare that it *specializes* the second bean:

```
@Alternative @Specializes
public class MockCreditCardPaymentProcessor
    extends CreditCardPaymentProcessor {
    ...
}
```

When an enabled bean specializes another bean, the other bean is never instantiated or called by the container. Even if the other bean defines a producer or observer method, the method will never be called.

So why does specialization work, and what does it have to do with inheritance?

Since we're informing the container that our alternative bean is meant to stand in as a replacement for the default implementation, the alternative implementation automatically inherits all qualifiers of the default implementation. Thus, in our example, `MockCreditCardPaymentProcessor` inherits the qualifiers `@Default` and `@CreditCard`.

Furthermore, if the default implementation declares a bean EL name using `@Named`, the name is inherited by the specialized alternative bean.

Java EE component environment resources

Java EE 5 already introduced some limited support for dependency injection, in the form of component environment injection. A component environment resource is a Java EE component, for example a JDBC datasource, JMS queue or topic, JPA persistence context, remote EJB or web service.

Naturally, there is now a slight mismatch with the new style of dependency injection in CDI. Most notably, component environment injection relies on string-based names to qualify ambiguous types, and there is no real consistency as to the nature of the names (sometimes a JNDI name, sometimes a persistence unit name, sometimes an EJB link, sometimes a non-portable "mapped name"). Producer fields turned out to be an elegant adaptor to reduce all this complexity to a common model and get component environment resources to participate in the CDI system just like any other kind of bean.

Fields have a duality in that they can both be the target of Java EE component environment injection and be declared as a CDI producer field. Therefore, they can define a mapping from a string-based name in the component environment, to a combination of type and qualifiers used in the world of typesafe injection. We call a producer field that represents a reference to an object in the Java EE component environment a *resource*.

14.1. Defining a resource

The CDI specification uses the term *resource* to refer, generically, to any of the following kinds of object which might be available in the Java EE component environment:

- JDBC Datasources, JMS Queues, Topics and ConnectionFactorys, JavaMail Sessions and other transactional resources including JCA connectors,
- JPA EntityManagers and EntityManagerFactorys,
- remote EJBs, and
- web services

We declare a resource by annotating a producer field with a component environment injection annotation: `@Resource`, `@EJB`, `@PersistenceContext`, `@PersistenceUnit` or `@WebServiceRef`.

```
@Produces @WebServiceRef(lookup="java:app/service/Catalog")
Catalog catalog;
```

```
@Produces @Resource(lookup="java:global/env/jdbc/CustomerDatasource")
@CustomerDatabase Datasource customerDatabase;
```

```
@Produces @PersistenceContext(unitName="CustomerDatabase")
@CustomerDatabase EntityManager customerDatabasePersistenceContext;
```

```
@Produces @PersistenceUnit(unitName="CustomerDatabase")
```

```
@CustomerDatabase EntityManagerFactory customerDatabasePersistenceUnit;
```

```
@Produces @EJB(ejbLink="../their.jar#PaymentService")  
PaymentService paymentService;
```

The field may be static (but not final).

A resource declaration really contains two pieces of information:

- the JNDI name, EJB link, persistence unit name, or other metadata needed to obtain a reference to the resource from the component environment, and
- the type and qualifiers that we will use to inject the reference into our beans.



Note

It might feel strange to be declaring resources in Java code. Isn't this stuff that might be deployment-specific? Certainly, and that's why it makes sense to declare your resources in a class annotated `@Alternative`.

14.2. Typesafe resource injection

These resources can now be injected in the usual way.

```
@Inject Catalog catalog;
```

```
@Inject @CustomerDatabase DataSource customerDatabase;
```

```
@Inject @CustomerDatabase EntityManager customerDatabaseEntityManager;
```

```
@Inject @CustomerDatabase EntityManagerFactory customerDatabaseEntityManagerFactory;
```

```
@Inject PaymentService paymentService;
```

The bean type and qualifiers of the resource are determined by the producer field declaration.

It might seem like a pain to have to write these extra producer field declarations, just to gain an additional level of indirection. You could just as well use component environment injection directly, right? But remember that you're going to be using resources like the `EntityManager` in several different beans. Isn't it nicer and more typesafe to write

```
@Inject @CustomerDatabase EntityManager
```

instead of

```
@PersistenceContext(unitName="CustomerDatabase") EntityManager
```

all over the place?

Part IV. CDI and the Java EE ecosystem

The third theme of CDI is *integration*. We've already seen how CDI helps integrate EJB and JSF, allowing EJBs to be bound directly to JSF pages. That's just the beginning. The CDI services are integrated into the very core of the Java EE platform. Even EJB session beans can take advantage of the dependency injection, event bus, and contextual lifecycle management that CDI provides.

CDI is also designed to work in concert with technologies outside of the platform by providing integration points into the Java EE platform via an SPI. This SPI positions CDI as the foundation for a new ecosystem of *portable* extensions and integration with existing frameworks and technologies. The CDI services will be able to reach a diverse collection of technologies, such as business process management (BPM) engines, existing web frameworks and de facto standard component models. Of course, 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 CDI 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 CDI. We'll also briefly meet a set of SPIs that are provided to support portable extensions to CDI. You might not ever need to use these SPIs directly, but don't take them for granted. You will likely be using them indirectly, every time you use a third-party extension, such as Seam.

Java EE integration

CDI is fully integrated into the Java EE environment. 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 other platform components, such as servlets and message-driven Beans, which are not beans themselves.

15.1. Built-in beans

In the Java EE environment, the container provides the following built-in beans, all with the qualifier `@Default`:

- the current JTA `UserTransaction`,
- a `Principal` representing the current caller identity,
- the default [Bean Validation](http://jcp.org/en/jsr/detail?id=303) [http://jcp.org/en/jsr/detail?id=303] `ValidationFactory`, and
- a `Validator` for the default `ValidationFactory`.



Note

The CDI specification does not require the servlet context objects, `HttpServletRequest`, `HttpSession` and `ServletContext` to be exposed as injectable beans. If you really want to be able to inject these objects, it's easy to create a portable extension to expose them as beans. However, we recommend that direct access to these objects be limited to servlets, servlet filters and servlet event listeners, where they may be obtained in the usual way as defined by the Java Servlets spec. The `FacesContext` is also not injectable. You can get at it by calling `FacesContext.getCurrentInstance()`.



Note

Oh, you *really* want to inject the `FacesContext`? Alright then, try this producer method:

```
class FacesContextProducer {
    @Produces @RequestScoped FacesContext getFacesContext() {
        return FacesContext.getCurrentInstance();
    }
}
```

15.2. Injecting Java EE resources into a bean

All managed beans may take advantage of Java EE component environment injection using `@Resource`, `@EJB`, `@PersistenceContext`, `@PersistenceUnit` and `@WebServiceRef`. 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 UserTransaction transaction;
```

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

```
@SessionScoped
public class Login implements Serializable {
    @Inject Credentials credentials;
    @PersistenceContext EntityManager userDatabase;
    ...
}
```

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

Of course, we advise that component environment injection be used to define CDI resources, and that typesafe injection be used in application code.

15.3. Calling a bean from a servlet

It's easy to use a bean from a servlet in Java EE 6. Simply inject the bean using field or initializer method injection.

```
public class Login extends HttpServlet {
    @Inject Credentials credentials;
    @Inject Login login;

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

Since instances of servlets are shared across all incoming threads, the bean 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.

15.4. Calling a bean from a message-driven bean

CDI injection applies to all EJBs, even when they aren't managed beans. In particular, you can use CDI injection in message-driven beans, which are by nature not contextual objects.

You can even use CDI interceptor bindings for message-driven Beans.


```

@Transactional @MessageDriven
public class ProcessOrder implements MessageListener {
    @Inject Inventory inventory;
    @PersistenceContext EntityManager em;

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

```

Please note that there is no session or conversation context available when a message is delivered to a message-driven bean. Only `@RequestScoped` and `@ApplicationScoped` beans are available.

But how about beans which *send* JMS messages?

15.5. 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.

You can use producer fields and methods to prepare all of these resources for injection into a bean:

```

public class OrderResources {
    @Resource(name="jms/ConnectionFactory")
    private ConnectionFactory connectionFactory;

    @Resource(name="jms/OrderQueue")
    private Queue orderQueue;

    @Produces @OrderConnection
    public Connection createOrderConnection() throws JMSEException {
        return connectionFactory.createConnection();
    }

    public void closeOrderConnection(@Disposes @OrderConnection Connection connection)
        throws JMSEException {
        connection.close();
    }

    @Produces @OrderSession
    public Session createOrderSession(@OrderConnection Connection connection)
        throws JMSEException {
        return connection.createSession(true, Session.AUTO_ACKNOWLEDGE);
    }

    public void closeOrderSession(@Disposes @OrderSession Session session)
        throws JMSEException {
        session.close();
    }

    @Produces @OrderMessageProducer
    public MessageProducer createOrderMessageProducer(@OrderSession Session session)

```

```
        throws JMSEException {
            return session.createProducer(orderQueue);
        }

        public void closeOrderMessageProducer(@Disposes @OrderMessageProducer MessageProducer producer)
            throws JMSEException {
            producer.close();
        }
    }
}
```

In this example, we can just inject the prepared `MessageProducer`, `Connection` or `QueueSession`:

```
@Inject Order order;
@Inject @OrderMessageProducer MessageProducer producer;
@Inject @OrderSession QueueSession orderSession;

public void sendMessage() {
    MapMessage msg = orderSession.createMapMessage();
    msg.setLong("orderId", order.getId());
    ...
    producer.send(msg);
}
```

The lifecycle of the injected JMS objects is completely controlled by the container.

15.6. Packaging and deployment

CDI doesn't define any special deployment archive. You can package beans in jars, ejb jars or wars—any deployment location in the application classpath. However, the archive must be a "bean archive". That means each archive that contains beans *must* include a file named `beans.xml` in the `META-INF` directory of the classpath or `WEB-INF` directory of the web root (for war archives). The file may be empty. Beans deployed in archives that do not have a `beans.xml` file will not be available for use in the application.

In an embeddable EJB container, beans may be deployed in any location in which EJBs may be deployed. Again, each location must contain a `beans.xml` file.

Portable extensions

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

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

More formally, according to the spec:

A portable extension may integrate with the container by:

- Providing its own beans, interceptors and decorators to the container
- Injecting dependencies into its own objects using the dependency injection service
- Providing a context implementation for a custom scope
- Augmenting or overriding the annotation-based metadata with metadata from some other source

16.1. Creating an `Extension`

The first step in creating a portable extension is to write a class that implements `Extension`. This marker interface does not define any methods, but it's needed to satisfy the requirements of Java SE's service provider architecture.

```
class MyExtension implements Extension { ... }
```

Next, we need to register our extension as a service provider by creating a file named `META-INF/services/javax.enterprise.inject.spi.Extension`, which contains the name of our extension class:

```
org.mydomain.extension.MyExtension
```

An extension is not a bean, exactly, since it is instantiated by the container during the initialization process, before any beans or contexts exist. However, it can be injected into other beans once the initialization process is complete.

```
@Inject
MyBean(MyExtension myExtension) {
    myExtension.doSomething();
}
```

And, like beans, extensions can have observer methods. Usually, the observer methods observe *container lifecycle events*.

16.2. Container lifecycle events

During the initialization process, the container fires a series of events, including:

- BeforeBeanDiscovery
- ProcessAnnotatedType
- ProcessInjectionTarget and ProcessProducer
- ProcessBean and ProcessObserverMethod
- AfterBeanDiscovery
- AfterDeploymentValidation

Extensions may observe these events:

```
class MyExtension implements Extension {

    void beforeBeanDiscovery(@Observes BeforeBeanDiscovery bbd) {
        Logger.global.debug("beginning the scanning process");
    }

    <T> void processAnnotatedType(@Observes ProcessAnnotatedType<T> pat) {
        Logger.global.debug("scanning type: " + pat.getAnnotatedType().getJavaClass().getName());
    }

    void afterBeanDiscovery(@Observes AfterBeanDiscovery abd) {
        Logger.global.debug("finished the scanning process");
    }

}
```

In fact, the extension can do a lot more than just observe. The extension is permitted to modify the container's metamodel and more. Here's a very simple example:

```
class MyExtension implements Extension {

    <T> void processAnnotatedType(@Observes ProcessAnnotatedType<T> pat) {
        //tell the container to ignore the type if it is annotated @Ignore
        if ( pat.getAnnotatedType().isAnnotationPresent(Ignore.class) ) pat.veto();
    }

}
```

The observer method may inject a BeanManager

```
<T> void processAnnotatedType(@Observes ProcessAnnotatedType<T> pat, BeanManager beanManager) { ... }
```

16.3. The `BeanManager` object

The nerve center for extending CDI is the `BeanManager` object. The `BeanManager` interface lets us obtain beans, interceptors, decorators, observers and contexts programmatically.

```
public interface BeanManager {
    public Object getReference(Bean<?> bean, Type beanType, CreationalContext<?> ctx);
    public Object getInjectableReference(InjectionPoint ij, CreationalContext<?> ctx);
    public <T> CreationalContext<T> createCreationalContext(Contextual<T> contextual);
    public Set<Bean<?>> getBeans(Type beanType, Annotation... qualifiers);
    public Set<Bean<?>> getBeans(String name);
    public Bean<?> getPassivationCapableBean(String id);
    public <X> Bean<? extends X> resolve(Set<Bean<? extends X>> beans);
    public void validate(InjectionPoint injectionPoint);
    public void fireEvent(Object event, Annotation... qualifiers);
    public <T> Set<ObserverMethod<? super T>> resolveObserverMethods(T event, Annotation... qualifiers);
    public List<Decorator<?>> resolveDecorators(Set<Type> types, Annotation... qualifiers);
    public List<Interceptor<?>> resolveInterceptors(InterceptionType type, Annotation... interceptorBindings);
    public boolean isScope(Class<? extends Annotation> annotationType);
    public boolean isNormalScope(Class<? extends Annotation> annotationType);
    public boolean isPassivatingScope(Class<? extends Annotation> annotationType);
    public boolean isQualifier(Class<? extends Annotation> annotationType);
    public boolean isInterceptorBinding(Class<? extends Annotation> annotationType);
    public boolean isStereotype(Class<? extends Annotation> annotationType);
    public Set<Annotation> getInterceptorBindingDefinition(Class<? extends Annotation> bindingType);
    public Set<Annotation> getStereotypeDefinition(Class<? extends Annotation> stereotype);
    public Context getContext(Class<? extends Annotation> scopeType);
    public ELResolver getELResolver();
    public ExpressionFactory wrapExpressionFactory(ExpressionFactory expressionFactory);
    public <T> AnnotatedType<T> createAnnotatedType(Class<T> type);
    public <T> InjectionTarget<T> createInjectionTarget(AnnotatedType<T> type);
}
```

Any bean or other Java EE component which supports injection can obtain an instance of `BeanManager` via injection:

```
@Inject BeanManager beanManager;
```

Java EE components may obtain an instance of `BeanManager` from JNDI by looking up the name `java:comp/BeanManager`. Any operation of `BeanManager` may be called at any time during the execution of the application.

Let's study some of the interfaces exposed by the `BeanManager`.

16.4. The `InjectionTarget` interface

The first thing that a framework developer is going to look for in the portable extension SPI is a way to inject CDI beans into objects which are not under the control of CDI. The `InjectionTarget` interface makes this very easy.



Note

We recommend that frameworks let CDI take over the job of actually instantiating the framework-controlled objects. That way, the framework-controlled objects can take advantage of constructor injection. However, if the framework requires use of a constructor with a special signature, the framework will need to instantiate the object itself, and so only method and field injection will be supported.

```
//get the BeanManager from JNDI
BeanManager beanManager = (BeanManager) new InitialContext().lookup("java:comp/BeanManager");

//CDI uses an AnnotatedType object to read the annotations of a class
AnnotatedType<SomeFrameworkComponent> type = beanManager.createAnnotatedType(SomeFrameworkComponent.class);

//The extension uses an InjectionTarget to delegate instantiation, dependency injection
//and lifecycle callbacks to the CDI container
InjectionTarget<SomeFrameworkComponent> it = beanManager.createInjectionTarget(type);

//each instance needs its own CDI CreationalContext
CreationalContext ctx = beanManager.createCreationalContext(null);

//instantiate the framework component and inject its dependencies
SomeFrameworkComponent instance = it.produce(ctx); //call the constructor
it.inject(instance, ctx); //call initializer methods and perform field injection
it.postConstruct(instance); //call the @PostConstruct method

...

//destroy the framework component instance and clean up dependent objects
it.preDestroy(instance); //call the @PreDestroy method
it.dispose(instance); //it is now safe to discard the instance
ctx.release(); //clean up dependent objects
```

16.5. The `Bean` interface

Instances of the interface `Bean` represent beans. There is an instance of `Bean` registered with the `BeanManager` object for every bean in the application. There are even `Bean` objects representing interceptors, decorators and producer methods.

The `Bean` interface exposes all the interesting things we discussed in [Section 2.1, “The anatomy of a bean”](#).

```
public interface Bean<T> extends Contextual<T> {
    public Set<Type> getTypes();
    public Set<Annotation> getQualifiers();
    public Class<? extends Annotation> getScope();
    public String getName();
    public Set<Class<? extends Annotation>> getStereotypes();
    public Class<?> getBeanClass();
    public boolean isAlternative();
    public boolean isNullable();
    public Set<InjectionPoint> getInjectionPoints();
}
```

```
}
```

There's an easy way to find out what beans exist in the application:

```
Set<Bean<?>> allBeans = beanManager.getBeans(Object.class, new AnnotationLiteral<Any>() {});
```

The Bean interface makes it possible for a portable extension to provide support for new kinds of beans, beyond those defined by the CDI specification. For example, we could use the Bean interface to allow objects managed by another framework to be injected into beans.

16.6. Registering a Bean

The most common kind of CDI portable extension registers a bean (or beans) with the container.

In this example, we make a framework class, `SecurityManager` available for injection. To make things a bit more interesting, we're going to delegate back to the container's `InjectionTarget` to perform instantiation and injection upon the `SecurityManager` instance.

```
public class SecurityManagerExtension implements Extension {

    void afterBeanDiscovery(@Observes AfterBeanDiscovery abd, BeanManager bm) {

        //use this to read annotations of the class
        AnnotatedType<SecurityManager> at = bm.createAnnotatedType(SecurityManager.class);

        //use this to instantiate the class and inject dependencies
        final InjectionTarget<SecurityManager> it = bm.createInjectionTarget(at);

        abd.addBean( new Bean<SecurityManager>() {

            @Override
            public Class<?> getBeanClass() {
                return SecurityManager.class;
            }

            @Override
            public Set<InjectionPoint> getInjectionPoints() {
                return it.getInjectionPoints();
            }

            @Override
            public String getName() {
                return "securityManager";
            }

            @Override
            public Set<Annotation> getQualifiers() {
                Set<Annotation> qualifiers = new HashSet<Annotation>();
                qualifiers.add( new AnnotationLiteral<Default>() {} );
                qualifiers.add( new AnnotationLiteral<Any>() {} );
                return qualifiers;
            }

        })
    }
}
```

```
@Override
public Class<? extends Annotation> getScope() {
    return SessionScoped.class;
}

@Override
public Set<Class<? extends Annotation>> getStereotypes() {
    return Collections.emptySet();
}

@Override
public Set<Type> getTypes() {
    Set<Type> types = new HashSet<Type>();
    types.add(SecurityManager.class);
    types.add(Object.class);
    return types;
}

@Override
public boolean isAlternative() {
    return false;
}

@Override
public boolean isNullable() {
    return false;
}

@Override
public SecurityManager create(CreationalContext<SecurityManager> ctx) {
    SecurityManager instance = it.produce(ctx);
    it.inject(instance, ctx);
    it.postConstruct(instance);
    return instance;
}

@Override
public void destroy(SecurityManager instance,
    CreationalContext<SecurityManager> ctx) {
    it.preDestroy(instance);
    it.dispose(instance);
    ctx.release();
}

} );
}
```

But a portable extension can also mess with beans that are discovered automatically by the container.

16.7. Wrapping an `AnnotatedType`

One of the most interesting things that an extension class can do is process the annotations of a bean class *before* the container builds its metamodel.

Let's start with an example of an extension that provides support for the use of `@Named` at the package level. The package-level name is used to qualify the EL names of all beans defined in that package. The portable extension uses the `ProcessAnnotatedType` event to wrap the `AnnotatedType` object and override the `value()` of the `@Named` annotation.

```
public class QualifiedNameExtension implements Extension {

    <X> void processAnnotatedType(@Observes ProcessAnnotatedType<X> pat) {

        //wrap this to override the annotations of the class
        final AnnotatedType<X> at = pat.getAnnotatedType();

        AnnotatedType<X> wrapped = new AnnotatedType<X>() {

            @Override
            public Set<AnnotatedConstructor<X>> getConstructors() {
                return at.getConstructors();
            }

            @Override
            public Set<AnnotatedField<? super X>> getFields() {
                return at.getFields();
            }

            @Override
            public Class<X> getJavaClass() {
                return at.getJavaClass();
            }

            @Override
            public Set<AnnotatedMethod<? super X>> getMethods() {
                return at.getMethods();
            }

            @Override
            public <T extends Annotation> T getAnnotation(final Class<T> annType) {
                if ( Named.class.equals(annType) ) {
                    class NamedLiteral
                        extends AnnotationLiteral<Named>
                        implements Named {
                        @Override
                        public String value() {
                            Package pkg = at.getClass().getPackage();
                            String unqualifiedName = at.getAnnotation(Named.class).value();
                            final String qualifiedName;
                            if ( pkg.isAnnotationPresent(Named.class) ) {
                                qualifiedName = pkg.getAnnotation(Named.class).value()
                                    + '.' + unqualifiedName;
                            }
                            else {
                                qualifiedName = unqualifiedName;
                            }
                            return qualifiedName;
                        }
                    }
                }
                return (T) new NamedLiteral();
            }
        }
    }
}
```

```
        else {
            return at.getAnnotation(annType);
        }
    }

    @Override
    public Set<Annotation> getAnnotations() {
        return at.getAnnotations();
    }

    @Override
    public Type getBaseType() {
        return at.getBaseType();
    }

    @Override
    public Set<Type> getTypeClosure() {
        return at.getTypeClosure();
    }

    @Override
    public boolean isAnnotationPresent(Class<? extends Annotation> annType) {
        return at.isAnnotationPresent(annType);
    }
};

pat.setAnnotatedType(wrapped);
}
```

Here's a second example, which adds the `@Alternative` annotation to any class which implements a certain Service interface.

```
class ServiceAlternativeExtension implements Extension {

    <T> void processAnnotatedType(@Observes ProcessAnnotatedType<T> pat) {

        final AnnotatedType<T> type = pat.getAnnotatedType();

        if ( Service.class.isAssignableFrom( type.getJavaClass() ) ) {

            //if the class implements Service, make it an @Alternative
            AnnotatedType<T> wrapped = new AnnotatedType<T>() {

                @Override
                public boolean isAnnotationPresent(Class<? extends Annotation> annotationType) {
                    return annotationType.equals(Alternative.class) ?
                        true : type.isAnnotationPresent(annotationType);
                }

                //remaining methods of AnnotatedType
                ...
            }
        }
    }
}
```

```

        pat.setAnnotatedType(wrapped);
    }
}

```

The `AnnotatedType` is not the only thing that can be wrapped by an extension.

16.8. Wrapping an `InjectionTarget`

The `InjectionTarget` interface exposes operations for producing and disposing an instance of a component, injecting its dependencies and invoking its lifecycle callbacks. A portable extension may wrap the `InjectionTarget` for any Java EE component that supports injection, allowing it to intercept any of these operations when they are invoked by the container.

Here's a CDI portable extension that reads values from properties files and configures fields of Java EE components, including servlets, EJBs, managed beans, interceptors and more. In this example, properties for a class such as `org.mydomain.blog.Blogger` go in a resource named `org/mydomain/blog/Blogger.properties`, and the name of a property must match the name of the field to be configured. So `Blogger.properties` could contain:

```

firstName=Gavin
lastName=King

```

The portable extension works by wrapping the container's `InjectionTarget` and setting field values from the `inject()` method.

```

public class ConfigExtension implements Extension {

    <X> void processInjectionTarget(@Observes ProcessInjectionTarget<X> pit) {

        //wrap this to intercept the component lifecycle
        final InjectionTarget<X> it = pit.getInjectionTarget();

        final Map<Field, Object> configuredValues = new HashMap<Field, Object>();

        //use this to read annotations of the class and its members
        AnnotatedType<X> at = pit.getAnnotatedType();

        //read the properties file
        String propsFileName = at.getClass().getSimpleName() + ".properties";
        InputStream stream = at.getJavaClass().getResourceAsStream(propsFileName);
        if (stream != null) {

            try {
                Properties props = new Properties();
                props.load(stream);
                for (Map.Entry<Object, Object> property : props.entrySet()) {
                    String fieldName = property.getKey().toString();
                    Object value = property.getValue();
                    try {
                        Field field = at.getJavaClass().getField(fieldName);
                        field.setAccessible(true);

```

```
        if ( field.getType().isAssignableFrom( value.getClass() ) ) {
            configuredValues.put(field, value);
        }
        else {
            //TODO: do type conversion automatically
            pit.addDefinitionError( new InjectionException(
                "field is not of type String: " + field ) );
        }
    }
    catch (NoSuchFieldException nsfe) {
        pit.addDefinitionError(nsfe);
    }
    finally {
        stream.close();
    }
}
}
catch (IOException ioe) {
    pit.addDefinitionError(ioe);
}
}

InjectionTarget<X> wrapped = new InjectionTarget<X>() {

    @Override
    public void inject(X instance, CreationalContext<X> ctx) {
        it.inject(instance, ctx);

        //set the values onto the new instance of the component
        for (Map.Entry<Field, Object> configuredValue: configuredValues.entrySet()) {
            try {
                configuredValue.getKey().set(instance, configuredValue.getValue());
            }
            catch (Exception e) {
                throw new InjectionException(e);
            }
        }
    }

    @Override
    public void postConstruct(X instance) {
        it.postConstruct(instance);
    }

    @Override
    public void preDestroy(X instance) {
        it.dispose(instance);
    }

    @Override
    public void dispose(X instance) {
        it.dispose(instance);
    }

    @Override
    public Set<InjectionPoint> getInjectionPoints() {
        return it.getInjectionPoints();
    }
}
```

```
@Override
public X produce(CreationalContext<X> ctx) {
    return it.produce(ctx);
}

};

pit.setInjectionTarget(wrapped);

}

}
```

There's a lot more to the portable extension SPI than what we've discussed here. Check out the CDI spec or Javadoc for more information. For now, we'll just mention one more extension point.

16.9. The `Context` interface

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

```
public interface Context {
    public Class<? extends Annotation> getScope();
    public <T> T get(Contextual<T> contextual, CreationalContext<T> creationalContext);
    public <T> T get(Contextual<T> contextual);
    boolean isActive();
}
```

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

Next steps

Because CDI is so new, there's not yet a lot of information available online. That will change over time. Regardless, the CDI specification remains the authority for information on CDI. The spec is less than 100 pages and is quite readable (don't worry, it's not like your Blu-ray player manual). Of course, it covers many details we've skipped over here. The spec is available on the [JSR-299 page](http://jcp.org/en/jsr/detail?id=299) [http://jcp.org/en/jsr/detail?id=299] at the JCP website.

The CDI reference implementation, Weld, is being developed at the [Seam project](http://seamframework.org/Weld) [http://seamframework.org/Weld]. The RI development team and the CDI spec lead blog at in.relation.to [http://in.relation.to]. This guide was originally based on a series of blog entries published there while the specification was being developed. It's probably the best source of information about the future of CDI, Weld and Seam.

We encourage you to follow the [weld-dev](https://lists.jboss.org/mailman/listinfo/weld-dev) [https://lists.jboss.org/mailman/listinfo/weld-dev] mailing list and to get involved in [development](http://seamframework.org/Weld/Development) [http://seamframework.org/Weld/Development]. If you are reading this guide, you likely have something to offer.

Part V. Weld Reference Guide

Weld is the reference implementation of JSR-299, and is used by JBoss AS and GlassFish to provide CDI services for Java Enterprise Edition (Java EE) applications. Weld also goes beyond the environments and APIs defined by the JSR-299 specification by providing support for a number of other environments (such as a servlet container such as Tomcat, or Java SE).

You might also want to check out the Weld Extensions project which provides portable extensions to CDI (such as log injection and additional annotations for creating beans), and Seam, which provides integrations with other view layers (such as GWT and Wicket), other frameworks (such as Drools) as well as extensions to the ecosystem (such as support for security).

If you want to get started quickly using Weld (and, in turn, CDI) with JBoss AS, GlassFish or Tomcat and experiment with one of the examples, take a look at [Chapter 6, Getting started with Weld](#). Otherwise read on for a exhaustive discussion of using Weld in all the environments and application servers it supports and the Weld extensions.

Application servers and environments supported by Weld

18.1. Using Weld with JBoss AS

If you are using JBoss AS 6.0, no additional configuration is required to use Weld (or CDI for that matter). All you need to do is make your application a bean archive by adding `META-INF/beans.xml` to the classpath or `WEB-INF/beans.xml` to the web root!



Note

Additionally, Weld Servlet supports JBoss EAP 5.1, to do this use the `jboss5` variant of Weld Servlet.

18.2. GlassFish

Weld is also built into GlassFish from V3 onwards. Since GlassFish V3 is the Java EE 6 reference implementation, it must support all features of CDI. What better way for GlassFish to support these features than to use Weld, the JSR-299 reference implementation? Just package up your CDI application and deploy.

18.3. Servlet containers (such as Tomcat or Jetty)

While JSR-299 does not require support for servlet environments, Weld can be used in a servlet container, such as Tomcat 6.0 or Jetty 6.1.



Note

There is a major limitation to using a servlet container. Weld doesn't support deploying session beans, injection using `@EJB` or `@PersistenceContext`, or using transactional events in servlet containers. For enterprise features such as these, you should really be looking at a Java EE application server.

Weld can be used as a library in an web application that is deployed to a Servlet container. You should place `weld-servlet.jar` within the `WEB-INF/lib` directory relative to the web root. `weld-servlet.jar` is an "uber-jar", meaning it bundles all the bits of Weld and CDI required for running in a Servlet container, for your convenience. Alternatively, you can use its component jars. A list of transitive dependencies can be found in the `META-INF/DEPENDENCIES.txt` file inside the `weld-servlet.jar` artifact.

You also need to explicitly specify the servlet listener (used to boot Weld, and control its interaction with requests) in `WEB-INF/web.xml` in the web root:

```
<listener>
  <listener-class>org.jboss.weld.environment.servlet.Listener</listener-class>
</listener>
```

18.3.1. Tomcat

Tomcat has a read-only JNDI, so Weld can't automatically bind the BeanManager extension SPI. To bind the BeanManager into JNDI, you should populate `META-INF/context.xml` in the web root with the following contents:

```
<Context>
  <Resource name="BeanManager"
    auth="Container"
    type="javax.enterprise.inject.spi.BeanManager"
    factory="org.jboss.weld.resources.ManagerObjectFactory"/>
</Context>
```

and make it available to your deployment by adding this to the bottom of `web.xml`:

```
<resource-env-ref>
  <resource-env-ref-name>BeanManager</resource-env-ref-name>
  <resource-env-ref-type>
    javax.enterprise.inject.spi.BeanManager
  </resource-env-ref-type>
</resource-env-ref>
```

Tomcat only allows you to bind entries to `java:comp/env`, so the BeanManager will be available at `java:comp/env/BeanManager`

Weld also supports Servlet injection in Tomcat 6.

18.3.2. Jetty

Like Tomcat, Jetty has a read-only JNDI, so Weld can't automatically bind the BeanManager. To bind the BeanManager to JNDI in Jetty 6, you should populate `WEB-INF/jetty-env.xml` with the following contents:

```
<!DOCTYPE Configure PUBLIC "-//Mort Bay Consulting//DTD Configure//EN"
"http://jetty.mortbay.org/configure.dtd">
<Configure id="webAppCtx" class="org.mortbay.jetty.webapp.WebAppContext">
  <New id="BeanManager" class="org.mortbay.jetty.plus.naming.Resource">
    <Arg><Ref id="webAppCtx"/></Arg>
    <Arg>BeanManager</Arg>
    <Arg>
      <New class="javax.naming.Reference">
        <Arg>javax.enterprise.inject.spi.BeanManager</Arg>
        <Arg>org.jboss.weld.resources.ManagerObjectFactory</Arg>
        <Arg/>
      </New>
    </Arg>
  </New>
</Configure>
```

Jetty 7 has moved to the Eclipse Foundation; if you are using Jetty 7 put the following in your `WEB-INF/jetty-env.xml`:

```

<!DOCTYPE Configure PUBLIC "-//Mort Bay Consulting//DTD Configure//EN"
"http://www.eclipse.org/jetty/configure.dtd">

<Configure id="webAppCtx" class="org.eclipse.jetty.webapp.WebAppContext">
  <New id="BeanManager" class="org.eclipse.jetty.plus.jndi.Resource">
    <Arg> <Ref id="webAppCtx"/> </Arg>
    <Arg>BeanManager</Arg>
    <Arg>
      <New class="javax.naming.Reference">
        <Arg>javax.enterprise.inject.spi.BeanManager</Arg>
        <Arg>org.jboss.weld.resources.ManagerObjectFactory</Arg>
        <Arg/>
      </New>
    </Arg>
  </New>
</Configure>

```

Just like in Tomcat, you need to make it available to your deployment by adding this to the bottom of `web.xml`:

```

<resource-env-ref>
  <resource-env-ref-name>BeanManager</resource-env-ref-name>
  <resource-env-ref-type>
    javax.enterprise.inject.spi.BeanManager
  </resource-env-ref-type>
</resource-env-ref>

```

Notice that Jetty doesn't not have built-in support for an `javax.naming.spi.ObjectFactory` like Tomcat, so it's necessary to manually create the `javax.naming.Reference` to wrap around it.

Jetty only allows you to bind entries to `java:comp/env`, so the `BeanManager` will be available at `java:comp/env/BeanManager`

Weld also supports Servlet injection in Jetty 6. To enable this, add the file `WEB-INF/jetty-web.xml` with the following content to your war:

```

<!DOCTYPE Configure PUBLIC "-//Mort Bay Consulting//DTD Configure//EN"
"http://jetty.mortbay.org/configure.dtd">
<Configure id="webAppCtx" class="org.mortbay.jetty.webapp.WebAppContext">
  <Call class="org.jboss.weld.environment.jetty.WeldServletHandler" name="process">
    <Arg><Ref id="webAppCtx"/></Arg>
  </Call>
</Configure>

```

18.4. Java SE

In addition to improved integration of the Enterprise Java stack, the "Contexts and Dependency Injection for the Java EE platform" specification also defines a state of the art typesafe, stateful dependency injection framework, which can prove useful in a wide range of application types. To help developers take advantage of this, Weld provides a simple means for being executed in the Java Standard Edition (SE) environment independently of any Java EE APIs.

When executing in the SE environment the following features of Weld are available:

- Managed beans with `@PostConstruct` and `@PreDestroy` lifecycle callbacks
- Dependency injection with qualifiers and alternatives
- `@Application`, `@Dependent` and `@Singleton` scopes
- Interceptors and decorators
- Stereotypes
- Events
- Portable extension support

EJB beans are not supported.

18.4.1. CDI SE Module

Weld provides an extension which will boot a CDI bean manager in Java SE, automatically registering all simple beans found on the classpath. The command line parameters can be injected using either of the following:

```
@Inject @Parameters List<String> params;
```

```
@Inject @Parameters String[] paramsArray;
```

The second form is useful for compatibility with existing classes.



Note

The command line parameters do not become available for injection until the `ContainerInitialized` event is fired. If you need access to the parameters during initialization you can do so via the public static `String[] getParameters()` method in `StartMain`.

Here's an example of a simple CDI SE application:

```
@Singleton
public class HelloWorld
{
    public void printHello(@Observes ContainerInitialized event, @Parameters List<String> parameters) {
        System.out.println("Hello " + parameters.get(0));
    }
}
```

18.4.2. Bootstrapping CDI SE

CDI SE applications can be bootstrapped in the following ways.

18.4.2.1. The ContainerInitialized Event

Thanks to the power of CDI's typesafe event model, application developers need not write any bootstrapping code. The Weld SE module comes with a built-in main method which will bootstrap CDI for you and then fire a `ContainerInitialized` event. The entry point for your application code would therefore be a simple bean which observes the `ContainerInitialized` event, as in the previous example.

In this case your application can be started by calling the provided main method like so:

```
java org.jboss.weld.environment.se.StartMain <args>
```

18.4.2.2. Programmatic Bootstrap API

For added flexibility, CDI SE also comes with a bootstrap API which can be called from within your application in order to initialize CDI and obtain references to your application's beans and events. The API consists of two classes: `Weld` and `WeldContainer`.

```
public class Weld
{
    /** Boots Weld and creates and returns a WeldContainer instance, through which
     * beans and events can be accessed. */
    public WeldContainer initialize() {...}

    /** Convenience method for shutting down the container. */
    public void shutdown() {...}
}
```

```
public class WeldContainer
{
    /** Provides access to all beans within the application. */
    public Instance<Object> instance() {...}

    /** Provides access to all events within the application. */
    public Event<Object> event() {...}

    /** Provides direct access to the BeanManager. */
    public BeanManager getBeanManager() {...}
}
```

Here's an example application main method which uses this API to initialize a bean of type `MyApplicationBean`.

```
public static void main(String[] args) {
    Weld weld = new Weld();
    WeldContainer container = weld.initialize();
```

```
container.select(MyApplicationBean.class).get();
weld.shutdown();
}
```

Alternatively the application could be started by firing a custom event which would then be observed by another simple bean. The following example fires `MyEvent` on startup.

```
public static void main(String[] args) {
    Weld weld = new Weld();
    WeldContainer container = weld.initialize();
    container.event().select(MyEvent.class).fire( new MyEvent() );
    weld.shutdown();
}
```

18.4.3. Thread Context

In contrast to Java EE applications, Java SE applications place no restrictions on developers regarding the creation and usage of threads. Therefore Weld SE provides a custom scope annotation, `@ThreadScoped`, and corresponding context implementation which can be used to bind bean instances to the current thread. It is intended to be used in scenarios where you might otherwise use `ThreadLocal`, and does in fact use `ThreadLocal` under the hood.

To use the `@ThreadScoped` annotation you need to enable the `RunnableDecorator` which 'listens' for all executions of `Runnable.run()` and decorates them by setting up the thread context beforehand, bound to the current thread, and destroying the context afterwards.

```
<beans>
  <decorators>
    <class>org.jboss.weld.environment.se.threading.RunnableDecorator</class>
  </decorator>
</beans>
```



Note

It is not necessary to use `@ThreadScoped` in all multithreaded applications. The thread context is not intended as a replacement for defining your own application-specific contexts. It is generally only useful in situations where you would otherwise have used `ThreadLocal` directly, which are typically rare.

18.4.4. Setting the Classpath

Weld SE comes packaged as a 'shaded' jar which includes the CDI API, Weld Core and all dependant classes bundled into a single jar. Therefore the only Weld jar you need on the classpath, in addition to your application's classes and dependant jars, is the Weld SE jar. If you are working with a pure Java SE application you launch using `java`, this may be simpler for you.

If you prefer to work with individual dependencies, then you can use the `weld-core` jar which just contains the Weld SE classes. Of course in this mode you will need to assemble the classpath yourself. This mode is useful, for example, if you wish to use an alternative slf4j log binding.

If you work with a dependency management solution such as Maven you can declare a dependency on `org.jboss.weld.se:weld-se-core`.

Context Management

19.1. Managing the built in contexts

Weld allows you to easily manage the built in contexts by injecting them and calling lifecycle methods. Weld defines two types of context, *managed* and *unmanaged*. Managed contexts can be activated (allowing bean instances to be retrieved from the context), invalidated (scheduling bean instances for destruction) and deactivated (stopping bean instances from being retrieved, and if the context has been invalidated, causing the bean instances to be destroyed). Unmanaged contexts are always active; some may offer the ability to destroy instances.

Managed contexts can either be *bound* or *unbound*. An unbound context is scoped to the thread in which it is activated (instances placed in the context in one thread are not visible in other threads), and is destroyed upon invalidation and deactivation. Bound contexts are attached to some external data store (such as the Http Session or a manually propagated map) by *associating* the data store with the context before calling activate, and dissociating the data store after calling deactivate.



Tip

Weld automatically controls context lifecycle in many scenarios such as HTTP requests, EJB remote invocations, and MDB invocations. Many of the extensions for CDI offer context lifecycle for other environments, it's worth checking to see if there is a suitable extension before deciding to manage the context yourself.

Weld provides a number of built in contexts, which are shown in [Table 19.1, “Available Contexts in Weld”](#).

Table 19.1. Available Contexts in Weld

Scope	Qualifiers	Context	Notes
@Dependent	@Default	DependentContext	The dependent context is unbound and unmanaged
@RequestScoped	@Unbound	RequestContext	An unbound request context, useful for testing
@RequestScoped	@Bound	RequestContext	A request context bound to a manually propagated map, useful for testing or non-Servlet environments
	@Default	BoundRequestContext	
@RequestScoped	@Http	RequestContext	A request context bound to a Servlet request, used for any Servlet based request context
	@Default	HttpRequestContext	
@RequestScoped	@Ejb	RequestContext	A request context bound to an interceptor's invocation context, used for EJB invocations outside of Servlet requests
	@Default	EjbRequestContext	
@ConversationScoped	@Bound	ConversationContext	A conversation context bound to two manually
	@Default	BoundConversationContext	

Scope	Qualifiers	Context	Notes
			propagated maps (one which represents the request and one which represents the session), useful for testing or non-Servlet environments
@ConversationScoped	@Http	ConversationContext	A conversation context bound to a Servlet request, used for any Servlet based conversation context
	@Default	HttpConversationContext	
@SessionScoped	@Bound	SessionContext	A session context bound to a manually propagated map, useful for testing or non-Servlet environments
	@Default	BoundSessionContext	
@SessionScoped	@Http	SessionContext	A session context bound to a Servlet request, used for any Servlet based session context
	@Default	HttpSessionContext	
@ApplicationScoped	@Default	ApplicationContext	An application context backed by an application scoped singleton, it is unmanaged and unbound but does offer an option to destroy all entries
@SingletonScoped	@Default	SingletonContext	A singleton context backed by an application scoped singleton, it is unmanaged and unbound but does offer an option to destroy all entries

Unmanaged contexts offer little of interest in a discussion about managing context lifecycles, so from here on in we will concentrate on the managed contexts (unmanaged contexts of course play a vital role in the functioning of your application and Weld!). As you can see from the table above, the managed contexts offer a number of different implementations for the same scope; in general, each flavor of context for a scope has the same API. We'll walk through a number of common lifecycle management scenarios below; armed with this knowledge, and the Javadoc, you should be able to work with any of the context implementations Weld offers.

We'll start simple with the `BoundRequestContext`, which you might use to provide the request scope outside of a Servlet request or EJB Invocation.

```

/* Inject the BoundRequestContext. */
/* Alternatively, you could look this up from the BeanManager */
@Inject BoundRequestContext requestContext;

...

/* Start the request, providing a data store which will last the lifetime of the request */
public void startRequest(Map<String, Object> requestDataStore) {

```

```

    // Associate the store with the context and activate the context
    requestContext.associate(requestDataStore);
    requestContext.activate();
}

/* End the request, providing the same data store as was used to start the request */
public void endRequest(Map<String, Object> requestDataStore) {
    try {
        /* Invalidate the request (all bean instances will be scheduled for destruction) */
        requestContext.invalidate();
        /* Deactivate the request, causing all bean instances to be destroyed (as the context
        is invalid) */
        requestContext.deactivate();
    } finally {
        /* Ensure that whatever happens we dissociate to prevent any memory leaks */
        requestContext.dissociate(requestDataStore);
    }
}
}

```

The bound session context works in much the same way, excepting that invalidating and deactivating the session context causes the any conversations in the session to be destroyed as well. The http session context and http request context also work similarly, and might be of use if you find yourself creating threads from an http request). The http session context additionally offers a method which can immediately destroy the context.



Note

Weld's session contexts are "lazy" and don't require a session to actually exist until a bean instance must be written.

The conversation context offers a few more options, which we will walk through here.

```

@Inject BoundConversationContext conversationContext;

...

/* Start a transient conversation */
/* Provide a data store which will last the lifetime of the request */
/* and one that will last the lifetime of the session */
public void startTransientConversation(Map<String, Object> requestDataStore,
                                     Map<String, Object> sessionDataStore) {
    resumeOrStartConversation(requestDataStore, sessionDataStore, null);
}

/* Start a transient conversation (if cid is null) or resume a non-transient */
/* conversation. Provide a data store which will last the lifetime of the request */
/* and one that will last the lifetime of the session */
public void resumeOrStartConversation(Map<String, Object> requestDataStore,
                                     Map<String, Object> sessionDataStore,
                                     String cid) {
    /* Associate the stores with the context and activate the context */
    /* BoundRequest just wraps the two datastores */
    conversationContext.associate(new MutableBoundRequest(requestDataStore, sessionDataStore));
    // Pass the cid in
    conversationContext.activate(cid);
}

```

```
}

/* End the conversations, providing the same data store as was used to start */
/* the request. Any transient conversations will be destroyed, any newly-promoted */
/* conversations will be placed into the session */
public void endOrPassivateConversation(Map<String, Object> requestDataStore,
                                       Map<String, Object> sessionDataStore) {

    try {
        /* Invalidate the conversation (all transient conversations will be scheduled for
destruction) */
        conversationContext.invalidate();
        /* Deactivate the conversation, causing all transient conversations to be destroyed */
        conversationContext.deactivate();
    } finally {
        /* Ensure that whatever happens we dissociate to prevent memory leaks*/
        conversationContext.dissociate(new MutableBoundRequest(requestDataStore, sessionDataStore));
    }
}
```

The conversation context also offers a number of properties which control the behavior of conversation expiration (after this period of inactivity the conversation will be ended and destroyed by the container), and the duration of lock timeouts (the conversation context ensures that a single thread is accessing any bean instances by locking access, if a lock can't be obtained after a certain time Weld will error rather than continue to wait for the lock). Additionally, you can alter the name of the parameter used to transfer the conversation id (by default, `cid`).

Weld also introduces the notion of a `ManagedConversation`, which extends the `Conversation` interface with the ability to lock, unlock and touch (update the last used timestamp) a conversation. Finally, all non-transient conversations in a session can be obtained from the conversation context, as can the current conversation.



Note

Weld's conversations are not assigned ids until they become non-transient.

Configuration

20.1. Excluding classes from scanning and deployment

Weld allows you to exclude classes in your archive from scanning, having container lifecycle events fired, and being deployed as beans.

In this tutorial, we'll explore this feature via an example; a more formal specification can be found in the xsd, located at http://jboss.org/schema/weld/beans_1_1.xsd.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:weld="http://jboss.org/schema/weld/beans"
  xsi:schemaLocation="
    http://java.sun.com/xml/ns/javaee http://docs.jboss.org/cdi/beans_1_0.xsd
    http://jboss.org/schema/weld/beans http://jboss.org/schema/weld/beans_1_1.xsd">

  <weld:scan>

    <!-- Don't deploy the classes for the swing app! -->
    <weld:exclude name="com.acme.swing.**" />

    <!-- Don't include GWT support if GWT is not installed -->
    <weld:exclude name="com.acme.gwt.**">
      <weld:if-class-available name="!com.google.GWT"/>
    </weld:exclude>

    <!--
      Exclude classes which end in Blether if the system property verbosity is set to low
      i.e.
      java ... -Dverbosity=low
    -->
    <weld:exclude pattern="^(.*)Blether$">
      <weld:if-system-property name="verbosity" value="low"/>
    </weld:exclude>

    <!--
      Don't include JSF support if Wicket classes are present, and the viewlayer system
      property is not set
    -->
    <weld:exclude name="com.acme.jsf.**">
      <weld:if-class-available name="org.apahce.wicket.Wicket"/>
      <weld:if-system-property name="!viewlayer"/>
    </weld:exclude>
  </weld:scan>

</beans>
```

In this example we show the most common use cases for exercising fine control over which classes Weld scans. The first filter excludes all classes in the package `com.acme.swing`, and in most cases this will be sufficient for your needs.

However, sometimes it's useful to be able to activate the filter depending on the environment used. In this case, Weld allows you to activate (or deactivate) a filter based on either system properties or whether a class is available. The second filter shows the use case of disabling scanning of certain classes depending on the capabilities of the environment you deploy to - in this case we are excluding GWT support if GWT is not installed.



Note

Notice how we use a `!` character on the `name` attribute to invert the activation condition. You can invert any activation condition in this way.

The third filter uses a regular expression to select the classes to match (normally we use simple name-based patterns, as they don't require us to escape the period which delimits the package hierarchy).



Note

If you specify just a system property name, Weld will activate the filter if that system property has been set (with any value). If you also specify the system property value, then Weld will only activate the filter if the system property's value matches exactly.

The fourth filter shows more advanced configurations, where we use multiple activation conditions to decide whether to activate the filter.

You can combine as many activation conditions as you like (*all* must be true for the filter to be activated). If you want a filter that is active if *any* of the activation conditions are true, then you need multiple identical filters, each with different activation conditions.



Tip

In general, the semantics defined by Ant's pattern sets (<http://ant.apache.org/manual/dirtasks.html#patternset>) are followed.

20.2. Parallel bootstrap configuration

Weld by default supports parallel loading and deploying of beans. However, in certain deployment scenarios the default setup may not be appropriate. Bootstrap configuration may be altered using the `org.jboss.weld.bootstrap.properties` file located on the classpath (e.g. `WEB-INF/classes/org.jboss.weld.bootstrap.properties` in a web archive).

Table 20.1. Bootstrap configuration options

Configuration option	Default value	Description
<code>enableThreading</code>	<code>true</code>	If set to <code>false</code> , Weld will only use the default thread for bootstrap and will not try to spawn any additional threads. This mode is safe for environments such as GAE.
<code>debug</code>	<code>false</code>	If set to <code>true</code> , debug timing information is printed to the standard output.

Configuration option	Default value	Description
deployerThreads	<code>Runtime.getRuntime().availableProcessors()</code>	The number of threads to be used for bean loading and deployment. If set to 0 the serial deployer is used instead of its concurrent counterpart.
preloaderThreads	<code>Math.max(1, Runtime.getRuntime().availableProcessors() - 1)</code>	Weld is capable of resolving observer methods (for container lifecycle events in advance while bean deployer threads are blocked waiting for I/O operations (such as classloading). This process is called preloading and leads to better CPU utilization and faster application startup time. This configuration option specifies the number of threads used for preloading. If set to 0, preloading is disabled.

Appendix A. Integrating Weld into other environments

If you want to use Weld in another environment, you will need to provide certain information to Weld via the integration SPI. In this Appendix we will briefly discuss the steps needed.



Enterprise Services

If you just want to use managed beans, and not take advantage of enterprise services (EE resource injection, CDI injection into EE component classes, transactional events, support for CDI services in EJBs) and non-flat deployments, then the generic servlet support provided by the "Weld: Servlets" extension will be sufficient, and will work in any container supporting the Servlet API.

All SPIs and APIs described have extensive JavaDoc, which spell out the detailed contract between the container and Weld.

A.1. The Weld SPI

The Weld SPI is located in the `weld-spi` module, and packaged as `weld-spi.jar`. Some SPIs are optional, and should only be implemented if you need to override the default behavior; others are required.

All interfaces in the SPI support the decorator pattern and provide a `Forwarding` class located in the `helpers` sub package. Additional, commonly used, utility classes, and standard implementations are also located in the `helpers` sub package.

Weld supports multiple environments. An environment is defined by an implementation of the `Environment` interface. A number of standard environments are built in, and described by the `Environments` enumeration. Different environments require different services to be present (for example a Servlet container doesn't require transaction, EJB or JPA services). By default an EE environment is assumed, but you can adjust the environment by calling `bootstrap.setEnvironment()`.

Weld uses a generic-typed service registry to allow services to be registered. All services implement the `Service` interface. The service registry allows services to be added and retrieved.

A.1.1. Deployment structure

An application is often comprised of a number of modules. For example, a Java EE deployment may contain a number of EJB modules (containing business logic) and war modules (containing the user interface). A container may enforce certain *accessibility* rules which limit the visibility of classes between modules. CDI allows these same rules to apply to bean and observer method resolution. As the accessibility rules vary between containers, Weld requires the container to *describe* the deployment structure, via the `Deployment` SPI.

The CDI specification discusses *Bean Deployment Archives* (BDAs)—archives which are marked as containing beans which should be deployed to the CDI container, and made available for injection and resolution. Weld reuses this description of *Bean Deployment Archives* in its deployment structure SPI. Each deployment exposes the BDAs which it contains; each BDA may also reference other which it can access. Together, the transitive closure of this graph forms the beans which are deployed in the application.

To describe the deployment structure to Weld, the container should provide an implementation of `Deployment`. `Deployment.getBeanDeploymentArchives()` allows Weld to discover the modules which make up

the application. The CDI specification also allows beans to be specified programmatically as part of the bean deployment. These beans may, or may not, be in an existing BDA. For this reason, Weld will call `Deployment.loadBeanDeploymentArchive(Class clazz)` for each programmatically described bean.

As programmatically described beans may result in additional BDAs being added to the graph, Weld will discover the BDA structure every time an unknown BDA is returned by `Deployment.loadBeanDeploymentArchive`.



Virtual BDAs

In a strict container, each BDA might have to explicitly specify which other BDAs it can access. However many containers will allow an easy mechanism to make BDAs bi-directionally accessible (such as a library directory). In this case, it is allowable (and reasonable) to describe all such archives as a single, 'virtual' `BeanDeploymentArchive`.

A container, might, for example, use a flat accessibility structure for the application. In this case, a single `BeanDeploymentArchive` would be attached to the `Deployment`.

`BeanDeploymentArchive` provides three methods which allow its contents to be discovered by Weld—`BeanDeploymentArchive.getBeanClasses()` must return all the classes in the BDA, `BeanDeploymentArchive.getBeansXml()` must return a data structure representing the `beans.xml` deployment descriptor for the archive, and `BeanDeploymentArchive.getEjbs()` must provide an EJB descriptor for every EJB in the BDA, or an empty list if it is not an EJB archive.

To aid container integrator, Weld provides a built-in `beans.xml` parser. To parse a `beans.xml` into the data-structure required by `BeanDeploymentArchive`, the container should call `Bootstrap.parseBeansXml(URL)`. Weld can also parse multiple `beans.xml` files, merging them to become a single data-structure. This can be achieved by calling `Bootstrap.parseBeansXml(Iterable<URL>)`.

BDA X may also reference another BDA Y whose beans can be resolved by, and injected into, any bean in BDA X. These are the accessible BDAs, and every BDA that is directly accessible by BDA X should be returned. A BDA will also have BDAs which are accessible transitively, and the transitive closure of the sub-graph of BDA X describes all the beans resolvable by BDA X.



Matching the classloader structure for the deployment

In practice, you can regard the deployment structure represented by `Deployment`, and the virtual BDA graph as a mirror of the classloader structure for a deployment. If a class can from BDA X can be loaded by another in BDA Y, it is accessible, and therefore BDA Y's accessible BDAs should include BDA X.

To specify the directly accessible BDAs, the container should provide an implementation of `BeanDeploymentArchive.getBeanDeploymentArchives()`.



Note

Weld allows the container to describe a circular graph, and will convert a graph to a tree as part of the deployment process.

Certain services are provided for the whole deployment, whilst some are provided per-BDA. BDA services are provided using `BeanDeploymentArchive.getServices()` and only apply to the BDA on which they are provided.

The contract for `Deployment` requires the container to specify the portable extensions (see chapter 12 of the CDI specification) which should be loaded by the application. To aid the container integrator, Weld provides the method `Bootstrap.loadExtensions(ClassLoader)` which will load the extensions for the specified classloader.

A.1.2. EJB descriptors

Weld delegates EJB 3 bean discovery to the container so that it doesn't duplicate the work done by the EJB container, and respects any vendor-extensions to the EJB definition.

The `EjbDescriptor` should return the relevant metadata as defined in the EJB specification. Each business interface of a session bean should be described using a `BusinessInterfaceDescriptor`.

A.1.3. EE resource injection and resolution services

All the EE resource services are per-BDA services, and may be provided using one of two methods. Which method to use is at the discretion of the integrator.

The integrator may choose to provide all EE resource injection services themselves, using another library or framework. In this case the integrator should use the EE environment, and implement the [Section A.1.8, “Injection Services”](#) SPI.

Alternatively, the integrator may choose to use CDI to provide EE resource injection. In this case, the `EE_INJECT` environment should be used, and the integrator should implement the [Section A.1.4, “EJB services”](#) [141], [Section A.1.7, “Resource Services”](#) and [Section A.1.5, “JPA services”](#).



Important

CDI only provides annotation-based EE resource injection; if you wish to provide deployment descriptor (e.g. `ejb-jar.xml`) injection, you must use [Section A.1.8, “Injection Services”](#).

If the container performs EE resource injection, the injected resources must be serializable. If EE resource injection is provided by Weld, the resolved resource must be serializable.



Tip

If you use a non-EE environment then you may implement any of the EE service SPIs, and Weld will provide the associated functionality. There is no need to implement those services you don't need!

A.1.4. EJB services

EJB services are split between two interfaces which are both per-BDA.

`EJBServices` is used to resolve local EJBs used to back session beans, and must always be provided in an EE environment. `EJBServices.resolveEjb(EjbDescriptor ejbDescriptor)` returns a wrapper—`SessionObjectReference`—around the EJB reference. This wrapper allows Weld to request a reference that implements the given business interface, and, in the case of SFSBs, both request the removal of the EJB from the container and query whether the EJB has been previously removed.

`EJBResolutionServices.resolveEjb(InjectionPoint ij)` allows the resolution of `@EJB` (for injection into managed beans). This service is not required if the implementation of [Section A.1.8, “Injection Services”](#) takes care of `@EJB` injection.

A.1.5. JPA services

Just as EJB resolution is delegated to the container, resolution of `@PersistenceContext` for injection into managed beans (with the `InjectionPoint` provided), is delegated to the container.

To allow JPA integration, the `JpaServices` interface should be implemented. This service is not required if the implementation of [Section A.1.8, “Injection Services”](#) takes care of `@PersistenceContext` injection.

A.1.6. Transaction Services

Weld delegates JTA activities to the container. The SPI provides a couple hooks to easily achieve this with the `TransactionServices` interface.

Any `javax.transaction.Synchronization` implementation may be passed to the `registerSynchronization()` method and the SPI implementation should immediately register the synchronization with the JTA transaction manager used for the EJBs.

To make it easier to determine whether or not a transaction is currently active for the requesting thread, the `isTransactionActive()` method can be used. The SPI implementation should query the same JTA transaction manager used for the EJBs.

A.1.7. Resource Services

The resolution of `@Resource` (for injection into managed beans) is delegated to the container. You must provide an implementation of `ResourceServices` which provides these operations. This service is not required if the implementation of [Section A.1.8, “Injection Services”](#) takes care of `@Resource` injection.

A.1.8. Injection Services

An integrator may wish to use `InjectionServices` to provide additional field or method injection over-and-above that provided by Weld. An integration into a Java EE environment may use `InjectionServices` to provide EE resource injection for managed beans.

`InjectionServices` provides a very simple contract, the `InjectionServices.aroundInject(InjectionContext ic);` interceptor will be called for every instance that CDI injects, whether it is a contextual instance, or a non-contextual instance injected by `InjectionTarget.inject()`.

The `InjectionContext` can be used to discover additional information about the injection being performed, including the target being injected. `ic.proceed()` should be called to perform CDI-style injection, and call initializer methods.

A.1.9. Security Services

In order to obtain the `Principal` representing the current caller identity, the container should provide an implementation of `SecurityServices`.

A.1.10. Bean Validation Services

In order to obtain the default `ValidatorFactory` for the application deployment, the container should provide an implementation of `ValidationServices`.

A.1.11. Identifying the BDA being addressed

When a client makes a request to an application which uses Weld, the request may be addressed at any of the BDAs in the application deployment. To allow Weld to correctly service the request, it needs to know which BDA the request is addressed at. Where possible, Weld will provide some context, but use of these by the integrator is optional.



Note

Most Servlet contains use a classloader-per-war, this may provide a good way to identify the BDA in use for web requests.

When Weld needs to identify the BDA, it will use one of these services, depending on what is servicing the request:

```
ServletServices.getBeanDeploymentArchive(ServletContext ctx)
    Identify the war in use. The ServletContext is provided for additional context.
```

A.1.12. The bean store

Weld uses a map like structure to store bean instances - `org.jboss.weld.context.api.BeanStore`. You may find `org.jboss.weld.context.api.helpers.ConcurrentHashMapBeanStore` useful.

A.1.13. The application context

Weld expects the Application Server or other container to provide the storage for each application's context. The `org.jboss.weld.context.api.BeanStore` should be implemented to provide an application scoped storage.

A.1.14. Initialization and shutdown

The `org.jboss.weld.bootstrap.api.Bootstrap` interface defines the initialization for Weld, bean deployment and bean validation. To boot Weld, you must create an instance of `org.jboss.weld.bootstrap.WeldBeansBootstrap` (which implements `Bootstrap`), tell it about the services in use, and then request the container start.

The bootstrap is split into phases, container initialization, bean deployment, bean validation and shutdown. Initialization will create a manager, and add the built-in contexts, and examine the deployment structure. Bean deployment will deploy any beans (defined using annotations, programmatically, or built in). Bean validation will validate all beans.

To initialize the container, you call `Bootstrap.startInitialization()`. Before calling `startInitialization()`, you must register any services required by the environment. You can do this by calling, for example, `bootstrap.getServices().add(JpaServices.class, new MyJpaServices())`. You must also provide the application context bean store.

Having called `startInitialization()`, the Manager for each BDA can be obtained by calling `Bootstrap.getManager(BeanDeploymentArchive bda)`.

To deploy the discovered beans, call `Bootstrap.deployBeans()`.

To validate the deployed beans, call `Bootstrap.validateBeans()`.

To place the container into a state where it can service requests, call `Bootstrap.endInitialization()`

To shutdown the container you call `Bootstrap.shutdown()`. This allows the container to perform any cleanup operations needed.

A.1.15. Resource loading

Weld needs to load classes and resources from the classpath at various times. By default, they are loaded from the Thread Context ClassLoader if available, if not the same classloader that was used to

load Weld, however this may not be correct for some environments. If this is case, you can implement `org.jboss.weld.spi.ResourceLoader`.

A.2. The contract with the container

There are a number of requirements that Weld places on the container for correct functioning that fall outside implementation of APIs.

Classloader isolation

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

Servlet

If you are integrating Weld into a Servlet environment you must register `org.jboss.weld.servlet.WeldListener` as a Servlet listener, either automatically, or through user configuration, for each CDI application which uses Servlet.

You must ensure that `WeldListener.contextInitialized()` is called after beans are deployed is complete (`Bootstrap.deployBeans()` has been called).

JSF

If you are integrating Weld into a JSF environment you must register `org.jboss.weld.jsf.WeldPhaseListener` as a phase listener.

If you are integrating Weld into a JSF environment you must register `org.jboss.weld.el.WeldELContextListener` as an EL Context listener.

If you are integrating Weld into a JSF environment you must register `org.jboss.weld.jsf.ConversationAwareViewHandler` as a delegating view handler.

If you are integrating Weld into a JSF environment you must obtain the bean manager for the module and then call `BeanManager.wrapExpressionFactory()`, passing `Application.getExpressionFactory()` as the argument. The wrapped expression factory must be used in all EL expression evaluations performed by JSF in this web application.

If you are integrating Weld into a JSF environment you must obtain the bean manager for the module and then call `BeanManager.getELResolver()`, The returned EL resolver should be registered with JSF for this web application.



Tip

There are a number of ways you can obtain the bean manager for the module. You could call `Bootstrap.getManager()`, passing in the BDA for this module. Alternatively, you could use the injection into Java EE component classes, or look up the bean manager in JNDI.

If you are integrating Weld into a JSF environment you must register `org.jboss.weld.servlet.ConversationPropagationFilter` as a Servlet listener, either automatically, or through user configuration, for each CDI application which uses JSF. This filter can be registered for all Servlet deployment safely.



Note

Weld only supports JSF 1.2 and above.

JSP

If you are integrating Weld into a JSP environment you must register `org.jboss.weld.el.WeldELContextListener` as an EL Context listener.

If you are integrating Weld into a JSP environment you must obtain the bean manager for the module and then call `BeanManager.wrapExpressionFactory()`, passing `Application.getExpressionFactory()` as the argument. The wrapped expression factory must be used in all EL expression evaluations performed by JSP.

If you are integrating Weld into a JSP environment you must obtain the bean manager for the module and then call `BeanManager.getELResolver()`, The returned EL resolver should be registered with JSP for this web application.



Tip

There are a number of ways you can obtain the bean manager for the module. You could call `Bootstrap.getManager()`, passing in the BDA for this module. Alternatively, you could use the injection into Java EE component classes, or look up the bean manager in JNDI.

Session Bean Interceptor

If you are integrating Weld into an EJB environment you must register the `aroundInvoke` method of `org.jboss.weld.ejb.SessionBeanInterceptor` as a EJB around-invoke interceptor for all EJBs in the application, either automatically, or through user configuration, for each CDI application which uses enterprise beans. If you are running in a EJB 3.1 environment, you should register this as an around-timeout interceptor as well.



Important

You must register the `SessionBeanInterceptor` as the inner most interceptor in the stack for all EJBs.

The weld-core.jar

Weld can reside on an isolated classloader, or on a shared classloader. If you choose to use an isolated classloader, the default `SingletonProvider`, `IsolatedStaticSingletonProvider`, can be used. If you choose to use a shared classloader, then you will need to choose another strategy.

You can provide your own implementation of `Singleton` and `SingletonProvider` and register it for use using `SingletonProvider.initialize(SingletonProvider provider)`.

Weld also provides an implementation of Thread Context Classloader per application strategy, via the `TCCLSingletonProvider`.

Binding the manager in JNDI

You should bind the bean manager for the bean deployment archive into JNDI at `java:comp/BeanManager`. The type should be `javax.enterprise.inject.spi.BeanManager`. To obtain the correct bean manager for the bean deployment archive, you may call `bootstrap.getBeanManager(beanDeploymentArchive)`

Performing CDI injection on Java EE component classes

The CDI specification requires the container to provide injection into non-contextual resources for all Java EE component classes. Weld delegates this responsibility to the container. This can be achieved using the CDI defined `InjectionTarget` SPI. Furthermore, you must perform this operation on the correct bean manager for the bean deployment archive containing the EE component class.

The CDI specification also requires that a `ProcessInjectionTarget` event is fired for every Java EE component class. Furthermore, if an observer calls `ProcessInjectionTarget.setInjectionTarget()` the container must use *the specified* injection target to perform injection.

To help the integrator, Weld provides `WeldManager.fireProcessInjectionTarget()` which returns the `InjectionTarget` to use.

```
// Fire ProcessInjectionTarget, returning the InjectionTarget
// to use
InjectionTarget it = weldBeanManager.fireProcessInjectionTarget(clazz);

// Per instance required, create the creational context
CreationalContext<?> cc = beanManager.createCreationalContext(null);

// Produce the instance, performing any constructor injection required
Object instance = it.produce();

// Perform injection and call initializers
it.inject(instance, cc);

// Call the post-construct callback
it.postConstruct(instance);

// Call the pre-destroy callback
it.preDestroy(instance);

// Clean up the instance
it.dispose();
cc.release();
```

The container may intersperse other operations between these calls. Further, the integrator may choose to implement any of these calls in another manner, assuming the contract is fulfilled.

When performing injections on EJBs you must use the Weld-defined SPI, `WeldManager`. Furthermore, you must perform this operation on the correct bean manager for the bean deployment archive containing the EJB.

```
// Obtain the EjbDescriptor for the EJB
// You may choose to use this utility method to get the descriptor
EjbDescriptor<?> ejbDescriptor = beanManager.getEjbDescriptor(ejbName);

// Get an the Bean object
Bean<?> bean = beanManager.getBean(ejbDescriptor);

// Create the injection target
InjectionTarget it = deploymentBeanManager.createInjectionTarget(ejbDescriptor);

// Per instance required, create the creational context
CreationalContext<?> cc = deploymentBeanManager.createCreationalContext(bean);

// Perform injection and call initializers
it.inject(instance, cc);

// You may choose to have CDI call the post construct and pre destroy
// lifecycle callbacks
```

```
// Call the post-construct callback
it.postConstruct(instance);

// Call the pre-destroy callback
it.preDestroy(instance);

// Clean up the instance
it.dispose();
cc.release();
```
