



CHAPTER 7

DEPENDENCY INJECTION




OBJECTIVES

After completing “Dependency Injection,” you will be able to:

- Understand the need for dependency injection in a container.
- Explain the inner workings of dependency injection in EJB containers, and recognize the impact of relying on different injection methods:
 - Injection by type vs. injection by name
 - Annotations vs. deployment descriptors
- Use dependency annotations to effect injection by type.
- Use the component environment to effect injection by name.
- Understand the need to use context lookups when interacting with stateful session beans.
- Use JNDI lookups to find remote beans.
- Use appropriate strategies to make injectable dependencies such as EJB references useful to all components in web applications.

Interdependent Systems

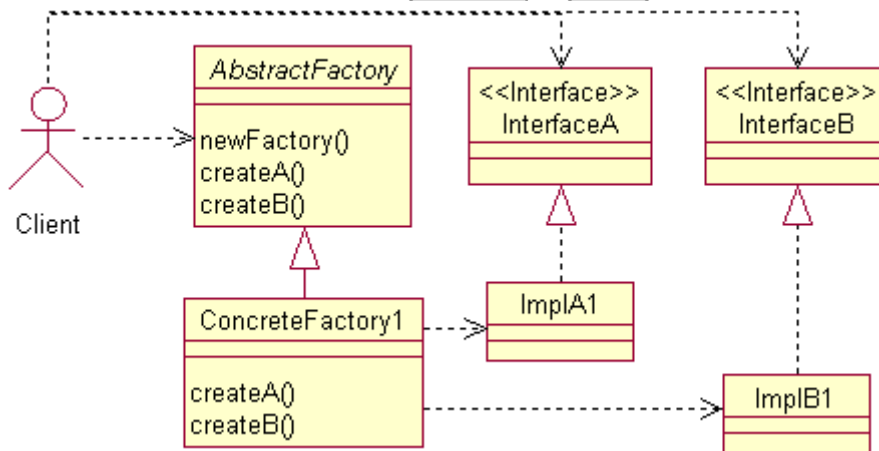
- Complex systems, by definition, function by interdependency; but at the same time dependencies bring maintenance challenges.
- A component (call it the **dependent**) that relies on another (call it the **collaborator**) for some functionality will best be defined to look to the interface of that other component.



- But as a practical matter, someone must choose a real implementation at runtime, and it's usually detrimental for the implementation to be chosen by the dependent component.
- This is **dependency on implementation type** and it makes the system less flexible, because alternate or replacement implementations might be referenced in many places.
- **Also, a dependent must find its collaborator somehow.**
 - In a single VM, this is a matter of knowing which **instance** of an object should be referenced.
 - In distributed systems, we're now talking about network-ready **names** for objects that can be published and discovered.
 - A **dependency on service location** has the same unfortunate effect on maintainability as a dependency on type: multiple points of maintenance that must be updated if the location changes.

The Factory Patterns

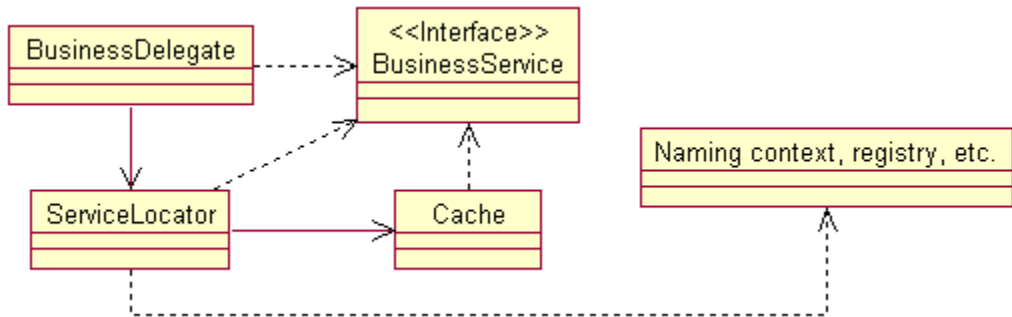
- In Java, one object often creates another object explicitly, as in
`HelperClass helper = new HelperClass ("module", 2);`
- Several closely-related design patterns define means of isolating the capability to create a certain object or family of objects – shown below is the **Abstract Factory** UML:



- Factories are the basic tools with which developers can get free of dependencies on implementation type.
- The key to the technique is that the body of code that is responsible for selecting the implementation type is neither the dependent nor the collaborator, but a third actor entirely.
 - The **factory method** pattern calls for a distinct method on one collaborator as a means of creating a different one – for instance a JDBC connection creating a statement and a statement creating a result set when running a query.

The Service Locator Pattern


- A Java EE design pattern, **Service Locator**, identifies a separate component with the responsibility to find required services for a dependent component:



- Service locators can do more than just locate services, as the diagram above suggests.
 - They can **cache** service locations.
 - They can **retry** and **restore** connections as needed.
 - They can **authenticate** the caller or the service.
- Service locators are a classic means of managing dependencies on name or location.

Dependency Injection

- EJB 3.0 – and Java EE 5.0 in general – defines the new feature of **dependency injection**:
 - A dependent declares its dependencies using Java annotations.
 - A Java-EE container resolves those dependencies by some heuristic involving class names, explicit bean names, and network locations.



- This provides a means of managing both types of dependency: dependency on implementation type and on name/location.

Injection by Magic

- Dependency injection at its simplest is a joy to use and may seem like outright magic:

- A bean class is declared `@Stateless`:

```
@Stateless  
public class MyClass  
    implements MyInterface
```

- ... and a client declares that it needs a bean of that type:

```
@EJB private MyClass service;
```

- And, presto! they start talking.

- There's more to the story, though, and for many common connectivity scenarios the actual usage is more complicated – though still well worth the effort.
- We'll start with the dead-simple usage as above, and work our way up and out to some more complicated cases.

- In **Examples/Wholesale/Step1** we see this “injection by magic” approach that we’ve been using for most of our EJB references so far.
 - An initialization servlet **cc.sales.web.InitServlet** publishes the injected reference to the **Records** service for all application components to use:

```
public class InitServlet
  extends HttpServlet
{
  public void init (ServletConfig config)
  {
    config.getServletContext ()
      .setAttribute ("records", service);
  }

  @EJB private Records service;
}
```

- **products.jsp** then uses this stored reference, as shown below. Other JSPs and servlets in the application all use this as their starting point for interacting with the order-fulfillment logic.

```
<c:forEach var="product"
  items="{records.allProducts}" >
```

Injection by Type

- Let's get back to that word we used earlier ... “magic.”
- Magic is great when it works, but unless we know the trick, we'll be lost when, for some reason, the magic stops working for us.
- So, how does dependency injection work?
- In the simple cases we've seen so far, the trick is just that the Java EE container will find a collaborating component of the correct **type** to satisfy the dependency.
 - The **@EJB** annotation modifies a field of a certain interface type.
 - The container seeks out an EJB that implements that type and, finding it, injects a reference to it to the annotated field.
- This suits a great many cases, because most often an EJB implements an interface that was designed just for it – or, to put it another way, there is usually only one class per interface.
- If there are multiple implementations of a bean interface, injection by type will not work.

```
@Stateless
public class Service1
    implements Service
...

```

```
@Stateless
public class Service2
    implements Service
...

```

- What then?

Injection by Name

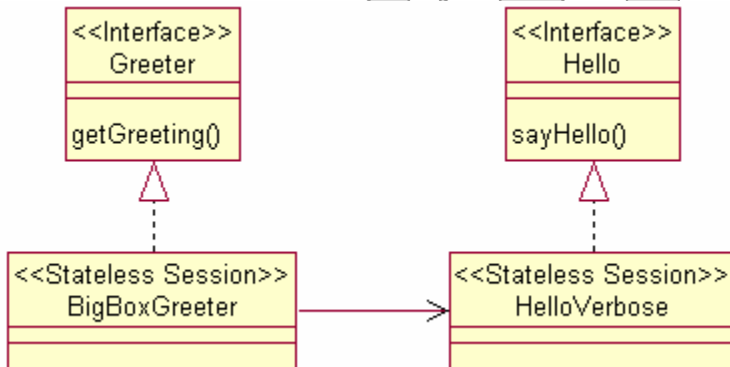
- To call out a specific collaborator, we use the name of a local EJB.
 - We'll come back to the remote case later in the chapter.
- The **@EJB** annotation defines a **beanName** attribute that can be used for this purpose.

```
@EJB (beanName="Service1")  
private Service service1;
```

```
@EJB (beanName="Service2")  
private Service service2;
```

- Don't confuse this with the **name** attribute.
- **name** provides the name of a bean when used in the **@Stateful** or **@Stateless** annotations; on the **@EJB** annotation it defines a name for the reference, not the name of a bean to use to resolve the dependency.

- Let's look at the differences between injection by type and injection by name.
- An expanded version of the “Hello, EJB” application includes two interfaces, each with a single implementation:



- We'll observe injection by type in the starter code, and then introduce a third bean and see how that breaks the system.
 - We'll then solve the problem with injection by name.
 - Do your work in **Demos/Injection**.
 - The completed demo is in **Examples/Hello/Step4**.
1. Review the code in `src/cc/ejb/BigBoxGreeter.java`, and see that it uses the simplest possible form of dependency injection to connect to **HelloVerbose** as a delegate:

```
public String getGreeting ()
{
    return delegate.sayHello () +
        " Can I help you find something today?";
}
```

```
@EJB private Hello delegate;
```

2. Build and test the new **Client** class, which now simply looks up and calls the **BigBoxGreeter**.

ant

runAC

Greeting from big-box greeter:

Well, hi, there! Nice to see you. Can I help you find something today?

3. We know that this works because the container looks for possible collaborators by the dependency type, and there is only one bean in the application that implements **Hello** right now.
4. Create a second bean class **HelloTerse** by copying **src/cc/ejb/HelloVerbose.java** and changing the class name and implementation:

```
@Stateless
public class HelloTerse
    implements Hello
{
    public String sayHello ()
    {
        return "Hi.";
    }
}
```

5. Build again. The deploy step may take quite a while this time, and when it fails, the error message says it all:

```
ant
```

```
...
```

```
Error resolving ejb-ref  
'cc.ejb.BigBoxGreeter/delegate' from module  
'Hello.jar' of application 'HelloApplication'. The  
ejb-ref does not have an ejb-link and the JNDI name  
of the target bean has not been specified. Attempts  
to automatically link the ejb-ref to its target  
bean failed because multiple EJBs in the  
application were found to implement the  
'cc.ejb.Hello' interface. Please specify a  
qualified ejb-link for this ejb-ref to indicate  
which EJB is the target of this ejb-ref.
```

6. Change **BigBoxGreeter** to select the bean it wants:

```
@EJB (beanName="HelloVerbose")  
private Hello delegate;
```

7. Rebuild and see that you get clean behavior once again.

```
ant
```

```
runAC
```

```
Greeting from big-box greeter:  
Well, hi, there! Nice to see you. Can I help you  
find something today?
```

Dependency “Injection?”

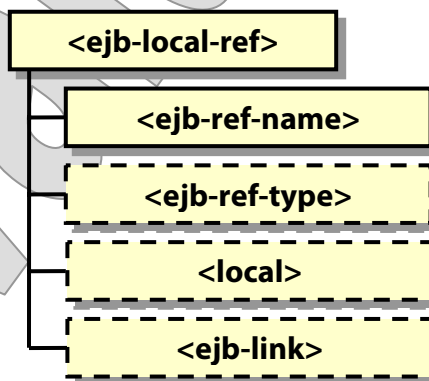
- When we use **beanName** to identify the collaborator, is this really dependency injection?
- Not exactly. Dependency injection implies that the collaborator is chosen by a separate actor (the container); in this usage the dependent source file identifies the collaborator explicitly.
- This is not dependency injection; it’s really just a slick way to look up an EJB by name.
- The practical impact is that, should the choice of bean change, we’d have to change the dependent source file and rebuild.
 - With proper dependency injection the dependent and collaborator components’ code is unchanged in this maintenance scenario.
 - Dependency by type meets this important criterion; using **beanName** does not.
 - This does not mean injection by name is a bad thing, but we’ll see next that there are other techniques to achieve it that are more flexible and maintainable than **beanName**.

The Component Environment

- Components can depend on things other than EJBs:
 - JDBC data sources
 - JMS message queues
 - Web services
 - Simple strings such as the URL of a required file
- For all these purposes, Java EE defines the concept of the **component environment**.
- This is essentially a private namespace in which useful resources can be registered.
- The mechanism is such that the dependent component can declare its environmental requirements to the container, and they can be satisfied in a variety of ways.
 - The component environment was essentially the earlier version of dependency injection for Java EE.
 - It's not as slick as Java-5 annotations, but it does essentially the same thing, which is to allow a component to connect to collaborators and resources without itself encoding dependencies on their types and locations.
 - In fact it has the advantage of being defined externally, avoiding even the naming dependency of the **beanName** technique.
- EJBs, servlets and other managed web components, and application clients all have component environments.

EJB Deployment Descriptors

- The deployment descriptor was an essential – and burdensome – element of EJB 2.1 development, defining all metadata.
- An EJB-3.0 application may have no deployment descriptor, since annotations can express almost all necessary metadata now.
- But the component environment transcends individual beans – by definition – so if environment entries are required, at least a minimal descriptor will be necessary.
- The subset of the deployment-descriptor model that governs component-environment definition of local EJB references is summarized here:



- At bottom this is a mapping from the `<ejb-ref-name>` to the `<ejb-link>`, which is the name of another bean.
- A similar model obtains for web and application-client components and their environment entries.

Finding Things

- The **EJBContext** interface offers the **lookup** method:
`public Object lookup (String localName);`
 - **SessionContext** and **EntityContext** extend **EJBContext**.
- This will look up an entry in the component environment by name, and return the entry as an object reference.
- We discussed development roles in an earlier chapter, and the component environment illustrates a clear distinction between two of those roles:
 - The **bean provider** chooses the name of the reference; this is appropriate because that name is used in the Java code for the component, typically as the argument to a **lookup** call.
 - The **application assembler** works at a scope that goes beyond a single bean, and is in a position to connect that private name to a resource unknown to the bean provider: another EJB, a data source, etc.
 - In this way each participant can carry out his or her tasks without stepping on the other's toes.
- The **name** attribute to the **@EJB** annotation is a shorthand for performing a context lookup during the initialization of the dependent component.

- Assume that we have an existing bean class in our application:

```
@Stateless public class CollaboratingBean
```

- Here is a hypothetical component environment for a session bean – this content would be found in a file **META-INF/ejb-jar.xml** located in the EJB JAR:

```
<ejb-jar version="3.0" ... >
  <enterprise-beans>
    <session>
      <ejb-name>DependentBean</ejb-name>
      <ejb-local-ref>
        <ejb-ref-name>myReference</ejb-ref-name>
        <ejb-ref-type>Session</ejb-ref-type>
        <local>com.you.Collaborator</local>
        <ejb-link>CollaboratingBean</ejb-link>
      </ejb-local-ref>
    </session>
  </enterprise-beans>
</ejb-jar>
```

- Code in the class **DependentBean** (or some other class annotated with a **name** of “DependentBean”) might perform a context lookup to connect to the collaborating bean:

```
private com.you.Collaborator otherBean;
@Resource private SessionContext context;

@PostConstruct public void init ()
{
  otherBean = (Collaborator)
    context.lookup ("myReference");
}
```

- Components in a web application share a component environment, defined at the top level in **web.xml**:

```
</web-app>
...
<ejb-local-ref>
  <ejb-ref-name>otherRef</ejb-ref-name>
  <ejb-ref-type>Session</ejb-ref-type>
  <local>com.you.Collaborator</local>
  <ejb-link>CollaboratingBean</ejb-link>
</ejb-local-ref>
</web-app>
```

- In code for a servlet somewhere in the application, we might take the simpler approach of defining a named, injectable reference to an EJB:

```
@EJB (name="otherRef") public Hello service;
```

- Again, note the important distinction between **name** and **beanName**:
 - **beanName** calls out the name of a collaborator explicitly.
 - **name** is the name of the reference itself, which belongs entirely to the bean in which the **@EJB** annotation is written; it is a port into which the collaborator can be injected via a lookup in the component environment.

- We'll do a quick cleanup of the Hello, EJB application in this regard – do your work in **Demos/CompEnv**.

- The completed demo is in **Examples/Hello/Step5**.

1. In **BigBoxGreeter**, change the **beanName** attribute to a **name** attribute and use this to define a reference called “delegate”.

```
@EJB (name="delegate") private Hello delegate;
```

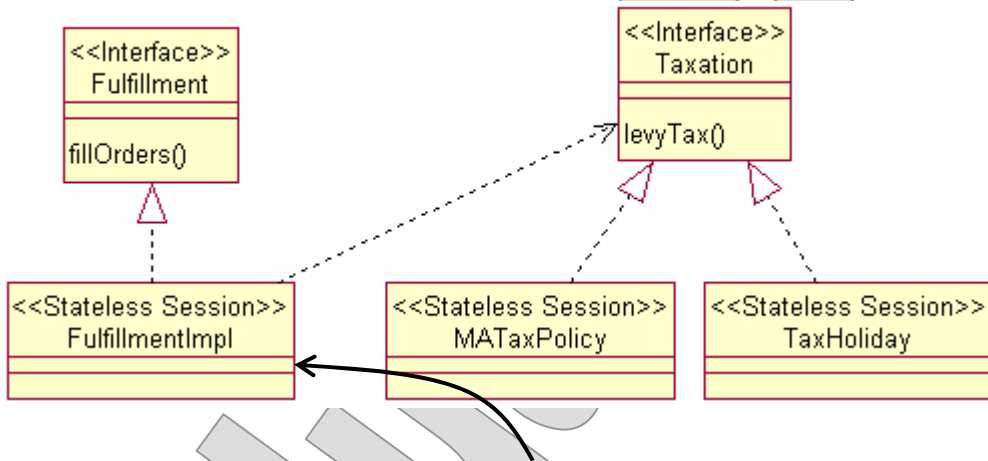
2. Copy the prepared deployment descriptor in **ejb-jar.xml** to the **ejb/META-INF** directory, and fill in the values as shown below. See what we're doing: we're declaring to the container that the **BigBoxGreeter** bean has a reference called **delegate** to a session bean with a business interface **cc.ejb.Hello**, and that this should be resolved to a bean named **HelloVerbose**.

```
<session>  
  <ejb-name>BigBoxGreeter</ejb-name>  
  <ejb-ref>  
    <ejb-ref-name>delegate</ejb-ref-name>  
    <ejb-ref-type>Session</ejb-ref-type>  
    <local>cc.ejb.Hello</local>  
    <ejb-link>HelloVerbose</ejb-link>  
  </ejb-ref>  
</session>
```

- Now, we have removed the information about the collaborating bean from the dependent bean, and define it externally. Changes in the wiring of the application are now manageable from this deployment descriptor, without making changes to (who knows how many) Java source files.

3. Build, deploy, and test, and see that the behavior is unchanged.

- In **Examples/Wholesale/Step2**, notice that the **Fulfillment** service used by the **ProcessingServlet** requires a pluggable tax policy – and that we use injection by name to achieve this:



– See `src/cc/sales/ejb/FulfillmentImpl.java`:

```
@EJB (name="taxPolicy") private Taxation taxPolicy;
```

– ... and `ejb/META-INF/ejb-jar.xml`:

```
<ejb-local-ref>  
  <ejb-ref-name>taxPolicy</ejb-ref-name>  
  <ejb-ref-type>Session</ejb-ref-type>  
  <local>cc.sales.ejb.Taxation</local>  
  <ejb-link>MATaxPolicy</ejb-link>  
</ejb-local-ref>
```

Impact on Stateful Session Beans

- Injection by type or injection by name – using **beanName**, reference **name**, or context lookups – works as advertised for collaborators that are stateful or stateless session beans.
- There is one additional consideration for stateful session beans, **if** the dependent component needs to create multiple instances of the stateful bean – perhaps for its own purposes or perhaps because it is acting as a factory for the stateful bean.
 - Recall that the container will automatically satisfy an **@EJB** dependency by creating a distinct stateful session bean – or at least by recognizing the reference as a distinct client of the bean, even if various beans from the pool are used to serve requests.
 - It will do the same for any explicit context **lookup**.
 - To create multiple instances on demand, though, an injectable dependency is unsuitable, because it is initialized just once, during the creation of the dependent component.
 - So context lookups, though they involve a few more lines of code, are the correct technique in such a case:

```
private MyService service;  
  
public void doGet (HttpServletRequest request,  
                  (HttpServletRequest response)  
                  throws ServletException, IOException  
{  
    request.getSession ().setAttribute ("bean",  
                                         aJNDIContext.lookup ("someStatefulBean");  
    ...  
}
```

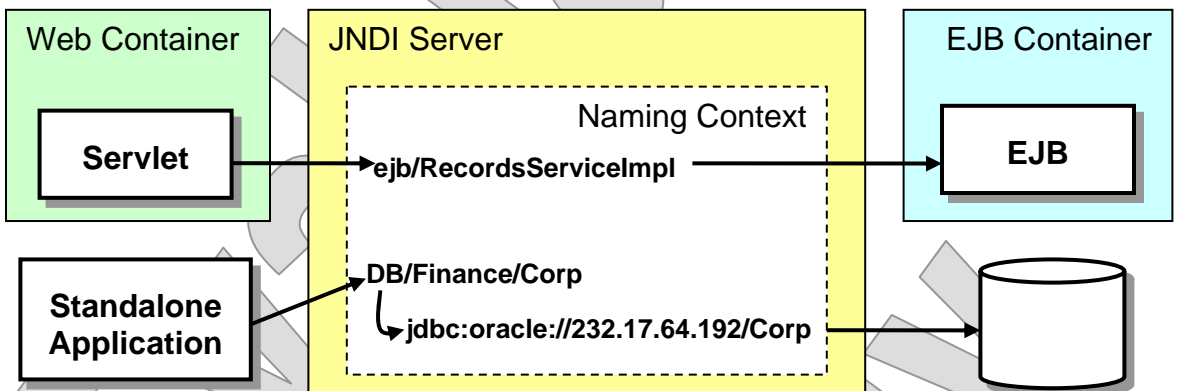
Suggested time: 15 minutes

In this lab you will connect or refactor `LandUse` and its bean connections, to fix the lingering problem that the stateless `LandUseServiceImpl` bean holds only one reference to the `ProposalServiceImpl` bean, and hands it out repeatedly, possibly to multiple callers who should not be sharing a stateful bean. You'll use a simple context lookup in the `getServiceForProposal` factory method; this will have the happy effect of prompting the EJB container to create a new instance of the stateful bean for each new lookup.

Detailed instructions are found at the end of the chapter.

Remote Connectivity and JNDI

- Simple **beanName** or `<ejb-link>` lookups are limited to local dependencies.
- A dependent must find a remote beans through other means – specifically, the **Java Naming and Directory Interface**, or **JNDI**.
- JNDI makes remote objects available on a network and allows Java clients to find them based on **compound names** – a fancy term that means JNDI manages a hierarchy of named nodes.



- An EJB container publishes its remote beans in the Java EE server's **global JNDI context**.
- The name under which a bean is published is determined one of two ways:
 - By a **non-portable deployment descriptor** that attaches a specific JNDI name to a specific bean
 - By the **mappedName** annotation attribute on the bean itself (which is also non-portable)

- In **Examples/JNDIReport** is a simple utility that does classic JNDI lookups using the **javax.naming** package.
 - It recurses through all published contexts and prints out a tree of names available on the local machine.
- **Build and test it now.** Depending on which code examples you've deployed, your exact results will vary:

ant

runAC

weblogic

HelloApplicationHello_jarBigBoxGreeter_Greeter

ejb

 mgmt

 MEJB

 BigBoxGreeter

 javax

 HelloApplicationHello_jarBigBoxGreeter_Home

 jdbc

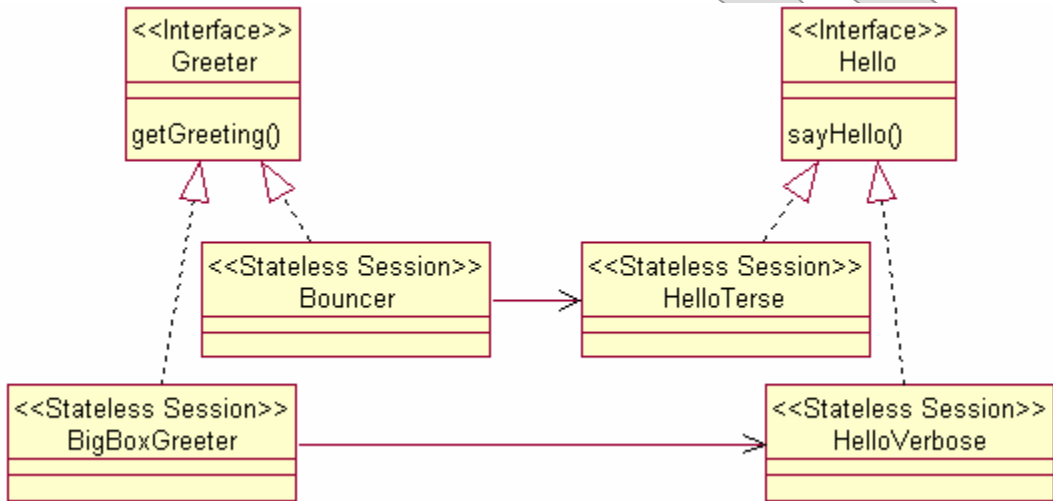
 LandUse

 Wholesale

 mejbmejb_jarMejb_EO

- **But notice the three names related to the **BigBoxGreeter** bean.**
 - The two more convoluted ones have to do with WebLogic publishing so-called home and remote objects to support EJB-2.x clients that may want to use this bean.
 - Then there's the name "ejb/BigBoxGreeter" ... how did that get there? We'll see in a moment.

- In **Examples/Hello/Step6**, we see an expanded set of beans, with a second type of **Greeter** that delegates to **HelloTerse**:



- We're also taking steps to set explicit JNDI names under which the two greeter beans will be published.
 - The new **Bouncer** bean uses the simple **mappedName** technique:
`@Stateless (mappedName="ejb/Bouncer")`
 - **BigBoxGreeter** takes the more formal approach, with no attributes on its **@Stateless** annotation but instead a non-portable deployment descriptor – see **ejb/META-INF/weblogic-ejb-jar.xml**:

```
<weblogic-enterprise-bean>
  <ejb-name>BigBoxGreeter</ejb-name>
  <stateless-session-descriptor>
    <business-interface-jndi-name-map>
      <business-remote>cc.ejb.Greeter</...>
      <jndi-name>ejb/BigBoxGreeter</jndi-name>
    </business-interface-jndi-name-map>
  ...
```

- Build and deploy the application and give it a try:

ant

runAC

```
Greeting from bouncer:  
Hi. ID, please?
```

Greeting from big-box greeter:

```
Well, hi, there! Nice to see you. Can I help you  
find something today?
```

- Run **JNDIReport** again and notice the different names that are now published – working from **Examples/JNDIReport**:

runAC

weblogic

```
HelloApplicationHello_jarBouncer_Home
```

```
HelloApplicationHello_jarBigBoxGreeter_Greeter
```

```
ejb
```

```
  mgmt
```

```
    MEJB
```

```
    BigBoxGreeter
```

```
    Bouncer#cc
```

```
      ejb
```

```
        Greeter
```

javax

```
HelloApplicationHello_jarBouncer_Greeter
```

```
HelloApplicationHello_jarBigBoxGreeter_Home
```

```
...
```

- Notice that while **BigBoxGreeter** is published at that clean, controlled name, **Bouncer** gets the **mappedName** – plus a suffix of “#cc/ejb/Greeter” added automatically by WebLogic, and based on the business interface implemented by the bean.

Finding Remote Things

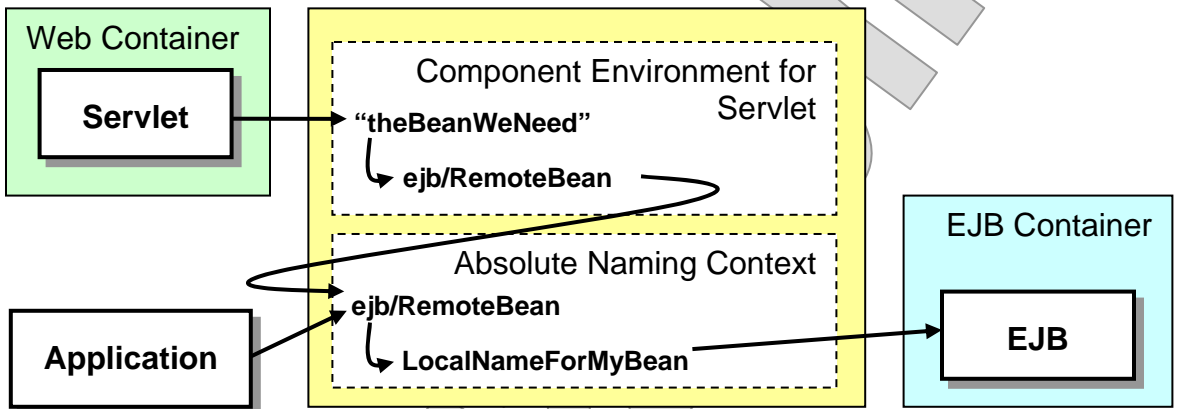
- A remote dependent must find its collaborating bean by performing a JNDI lookup.
- There are more and less complex ways in which to do this.
- The classic JNDI lookup incantation, used heavily in EJB 2.x and still possible for 3.0, is this:

```
Context ctx = new InitialContext ();  
BizInterface bean = (BizInterface)  
    javax.rmi.PortableRemoteObject.narrow  
        (ctx.lookup ("ejb/aBean"), BizInterface.class);
```

- This has two significant disadvantages:
 - The code is intricate and error-prone. Not shown is the required exception-handling around names not found, classes not found, and so on.
 - The JNDI name is encoded in the source file of the dependent component – again, not what we want from the standpoint of best maintainability.

JNDI for Dependency Injection

- EJB 3.0 offers an easier approach: references in the component environment can be wired to global JNDI names.



- This again requires the use of a **non-portable** deployment descriptor – now using the implicitly remote `<ejb-ref>`:

```
<ejb-ref>
  <ejb-ref-name>theBeanWeNeed</ejb-ref-name>
  <jndi-name>ejb/RemoteBean</jndi-name>
</ejb-ref>
```

- The dependent component can use **context lookups** the same way it does for local beans.

```
@Resource private SessionContext context;
...
context.lookup ("theBeanWeNeed");
```

- An important feature here is that the dependent is unaware not only of the name and location of the collaborator but even whether it is local or remote; this is known as **location transparency**.
- Again, an **@EJB** with a **name** attribute will be injected with the results of an automated context lookup.

Using mappedName for JNDI Lookup

- The **mappedName** attribute may be available on the **@EJB** annotation as well as on the bean-type annotations.
 - That is, support for this attribute is optional in the specification.
 - WebLogic does not support **mappedName** for resource or EJB references – only for EJBs themselves as a way of publishing them.
- This is **non-portable**, and **re-introduces the dependency on the global name of the collaborating object**.
 - In other words, while it is much simpler and easier to use, it suffers the same downsides as the EJB-2.1-style JNDI lookup we saw earlier.
 - It's fine for early development cycles or proofs of concept, but not robust enough for later development and production.

- We'll now review the client side of the Hello application more closely, and then tune the whole thing up a bit.
- Do your work in **Demos/Lookup**.
 - The completed demo is in **Examples/Hello/Step7**.
- First, look at the client application in **src/Client.java**. It uses EJB references initialized via dependency injection:

```
@EJB private static Greeter bouncer;  
@EJB private static Greeter bigBox;
```

- These requirements are met by the application client container (ACC) by looking at the non-portable deployment descriptor.
 - See **client/META-INF/weblogic-application-client.xml**:

```
<ejb-ref>  
  <ejb-ref-name>Client/bigBox</ejb-ref-name>  
  <jndi-name>ejb/BigBoxGreeter</jndi-name>  
</ejb-ref>  
<ejb-ref>  
  <ejb-ref-name>Client/bouncer</ejb-ref-name>  
  <jndi-name>ejb/Bouncer#cc.ejb.Greeter</jndi-name>  
</ejb-ref>
```

1. How are the `<ejb-ref-name>`s determined? By default, WebLogic will synthesize a name for the reference itself by combining the class and field names. But we can name them explicitly as well. In **Client.java**, add **name** attributes to the **@EJB** annotations:

```
@EJB(name="bouncer") private static Greeter bouncer;  
@EJB(name="bigBox") private static Greeter bigBox;
```

2. Now edit the deployment descriptor: just remove the “Client/” prefix from each reference name, since we’ve chosen to declare names that actually do match the field names that were already in play.

```
<ejb-ref>
  <ejb-ref-name>Client/bigBox</ejb-ref-name>
  <jndi-name>ejb/BigBoxGreeter</jndi-name>
</ejb-ref>
<ejb-ref>
  <ejb-ref-name>Client/bouncer</ejb-ref-name>
  <jndi-name>ejb/Bouncer#cc.ejb.Greeter</jndi-name>
</ejb-ref>
```

3. Make similar changes in **client/META-INF/application-client.xml**; this is the portable descriptor and must stay in synch as well.
4. You can build and test this refactored client code – note that you don’t need to redeploy the EJB application, so you can just build the client JAR again and run that:

```
ant build
```

```
runAC
```

```
Greeting from bouncer:
```

```
Hi. ID, please?
```

```
Greeting from big-box greeter:
```

```
Well, hi, there! Nice to see you. Can I help you
find something today?
```

5. Now let’s make a change to the EJB application. Open **src/cc/ejb/Bouncer.java**, and get rid of the **mappedName** attribute:

```
@Stateless(mappedName="ejb/Bouncer")
@Remote
public class Bouncer
```

- Open `ejb/META-INF/weblogic-ejb-jar.xml`. Clone the declaration of the **BigBoxGreeter** bean, and adjust it for **Bouncer** as follows:

```
<weblogic-enterprise-bean>
  <ejb-name>Bouncer</ejb-name>
  <stateless-session-descriptor>
    <business-interface-jndi-name-map>
      <business-remote>cc.ejb.Greeter
      </business-remote>
      <jndi-name>ejb/Bouncer</jndi-name>
    </business-interface-jndi-name-map>
  </stateless-session-descriptor>
</weblogic-enterprise-bean>
```

- Of course now we'll have to tweak the client so that it can find the relocated bean. This is a small illustration of the advantage of external deployment descriptors, in that we do not need to edit and recompile the application code – just the client-side XML:

```
<ejb-ref>
  <ejb-ref-name>Client/bouncer</ejb-ref-name>
  <jndi-name>ejb/Bouncer#cc.ejb.Greeter</jndi-name>
</ejb-ref>
```

- Rebuild – we do need a full undeploy-build-deploy cycle this time – and test once again:

ant

runAC

Greeting from bouncer:

Hi. ID, please?

Greeting from big-box greeter:

Well, hi, there! Nice to see you. Can I help you find something today?

Summary of Techniques

- Let's review the options available, as we've covered quite a bit of ground in this chapter.
- On the service side, an EJB can be named in several ways:
 - Its **default local name** is the class name with no package qualifiers.
 - We can override that with an explicit **name** attribute to the bean-type annotation (**@Stateful** or **@Stateless**).
 - It's **default remote name** is derived from the fully-qualified name of the business interface – not the class.
 - We can override this with the **mappedName** attribute.
 - Or, we can define it in a **non-portable deployment descriptor**.
- Clients can find local beans in a few ways:
 - **Injection by type**, if there is only one local implementation
 - Injection by **beanName** – this is suspect because it actually sets a name dependency into the source code
 - Injection by **portable deployment descriptor** and `<ejb-link>`, using either explicit **EJBContext.lookup** or a **name** for the dependency
- For remote beans the options are similar:
 - **Injection by type**, if there is only one implementation on the host
 - Injection by **mappedName** – on some servers, not including WebLogic, and this is a compromise in any case
 - Injection by **JNDI lookup** using a non-portable descriptor, and again either by explicit lookup or a reference **name**

Who Can Declare Dependencies

- So far we've focused on injection into a few types of components: EJBs, servlets, and applications running in the Java EE application-client container.
- Other objects can use dependency injection. EJB and resource references can be declared on the following component types:
 - EJBs
 - Servlets
 - Servlet filters
 - Web context, request, and session **event listeners**
 - Custom JSP **tag handlers**
 - JSF **managed beans**
 - **Application clients** in the Java EE container

EJBs in Web Applications

- Conspicuously absent are POJOs running in the web container.
- These classes, even though their code may be invoked directly from a servlet, have no direct relationship with the web container.
 - They are not declared to the container in **web.xml**, nor annotated as playing any special role.
 - Therefore the container cannot index them, manage them, control their lifecycle, inject dependencies, etc.
- For servlet/JSP web applications (or those that use MVC frameworks such as Struts and Spring), a strategy is needed to assure that objects that use EJBs and resources can find them.
- Here are three common approaches:
 - If a servlet is involved in each HTTP request, it can use direct dependency injection or JNDI lookup to find collaborators, and make either them or the results of invoking their methods available as **request attributes**.
 - When working with stateful session beans, a variant of this strategy has the servlet posting the stateful bean reference as a **session attribute**, which matches up nicely to the lifecycle and availability of the stateful bean itself.
 - An **initialization servlet** can publish its injected dependencies as **application attributes** through the servlet context.
- JSF applications have an easier time of it, since they have their own dependency-injection features and these integrate easily with EJB injection.

Suggested time: 15 minutes

In this lab you will observe the configuration and functioning of a simple web application deployed with the Hello/Greeter EJBs; and then configure a separate version of that web application that must function outside the main EAR deployment, contacting the target EJB remotely.

Detailed instructions are found at the end of the chapter.

SUMMARY

- **At first blush, dependency injection looks too good to be true.**
 - Too much of the literature on EJB3 boasts about the simplicity of injection by type, while conveniently ignoring the slight complications of injection by name.
- **In fact it is a powerful feature that is well worth the trouble to implement – but it is not the mind-reading system we’re sometimes led to expect.**
 - You can have the absolute simplicity of an unattributed **@EJB** annotation.
 - You can have the declarative power of choosing a collaborating bean by name (local or JNDI).
 - You just can’t have them both for the same dependency! and the cost of injection by name is not that high – nothing like the onerous EJB 2.1 descriptor that many developers remember with some dread.
- **When dealing with stateful session beans, it’s essentially unworkable to inject the EJB reference: instead, inject a context reference and look up the bean each time it’s needed.**
 - This will trigger creation of a new stateful bean for each new client conversation.