



Capstone Courseware, LLC

33 Boylston Street
Jamaica Plain, MA 02130

877-227-2477
capstonecourseware.com

The Spring Framework

Version 2.0.1

Instructor's Guide



Overview

Let's begin this guide for instructors more or less the way we begin the coursebook: by taking cognizance of the sheer size of the Spring framework. Certainly, no three-day course is going to do justice to the full scope and depth of this code library. This is by no means a survey course, and the boundaries around the course scope are more in the dimension of breadth than depth. But the boundaries are nevertheless there: we can't delve into every detail of the framework.

Also, with Spring as with any open-source tool, we're working with the brilliant ideas of a few bright people; we are not teaching an elected standard such as Servlets or JSP. And with a project as popular and dynamic as Spring, that step apart from standards-based frameworks is not the last one; there is a lively aftermarket in Spring utilities, helpers, IDEs, etc.

All of this suggests that it is better to teach this course from a position of enthusiasm for what the tool can do, than by projecting a sense of explaining exactly how it's supposed to be done. That is, if in Servlets training we say, "thou shalt extend HttpServlet," in Spring we tend to say, "here's one way to get your controllers mapped neatly to your requests." For every way of doing something in Spring that this course can show, there will be several others, and some of them will involve code not found in the Spring JARs. Spring is in a way the glue between the well-considered but somewhat static Java EE platform and the best practices and inspirations of any experienced and dedicated Java developer. The most responsive audiences will be those with some real Java EE experience under their belts – those who've done it the hard way a few times and will appreciate the subtle advantages of Spring's inversion of control and general respect for design patterns. Novice audiences – I'm thinking of the students who technically have the prerequisites but maybe not such deep experience – will probably need a little more in the way of evangelism: don't hesitate to get up on that soapbox if it seems necessary, because the details will be a lot more interesting to students who buy the argument in the first place that IoC is good, strategies are good, MVC is good, etc.





Timeline

The following breakdowns are approximate, and every class will vary.

Day 1

2 hours	Chapter 1
2½ hours	Chapter 2
2 hours, runs to day 2	Chapter 3

Day 2

1 hours	Chapter 3
4 hours	Chapter 4
1½ hours, runs to day 3	Chapter 5

Day 3

1 hour	Chapter 5
2½ hours	Chapter 6
1½ hours	Chapter 7
1½ hours	Chapter 8

All the demos in Chapters 4-6 can slow these down quite a bit. This is partly by design, as we really want to live in the space of the web module for a while, and not rush. But keep an eye on these, and if Chapter 4 seems to be getting you into time trouble on day 2, it's a good idea to ease up on the long interactive demos in the next couple chapters – or plan to compress or skip Chapter 8, which for some audiences may be just fine. The easiest way to get some time back is to take a demo as an example: that is, don't carry out the coding steps, but instead look in the directory called out as the “completed demo” and review the solution.

As mentioned in the outline, for less experienced Java developers, a four-day timeline is entirely workable. Some informal backgrounding on Java web applications with servlets and JSP is usually a good idea in this case, and that may occupy a couple hours; then simply spread out the above timeline by giving an extra 15-30 minutes to each lab exercise, letting students complete what they might not in the suggested lab times in the coursebook.





Crimson and the XML Validator

For those classrooms not using Eclipse, it will be worth the trouble to introduce students to a simple XML authoring environment deployed with the course software. The setup guide requires that the Crimson text editor be installed to `c:\Capstone\Tools\Crimson3.70`; then, the lab installer bundles an XML validating tool from Capstone Courseware (just a command-line interface to Xerces, really) and a single file that installs custom hotkeys into Crimson. The end result is that Crimson acts as an integrated editor, parser and schema validator for XML files. Hit **F9** to parse for well-formedness and **F10** to validate against a schema that's referenced in the XML document using the `xsi:schemaLocation` attribute.

Students may find this helpful when creating and editing their beans XML files and web context configurations. Especially if you sense that students' XML experience is slight, take a few minutes to show them this environment, and suggest that they make a habit of saving with **F10** rather than simply **File|Save**; this will trigger schema validation of their XML files and alert them to any of the common syntax or validity mistakes that can otherwise become frustrating as they don't surface until after the whole application's been run or deployed to Tomcat.

Two unfortunate limitations: one is that, on the one hand, Spring's bean factories insist on a `schemaLocation` attribute, or they fail to load the XML. This is non-compliant with XML Schema, but fairly harmless as we prime all our files with this attribute. On the other hand, web application contexts will choke if the `schemaLocation` is there and can't be resolved at runtime! So for safety's sake we leave them out of the context configurations in all the web projects. This in turn means that **F10** won't be useful for these files; students can still hit **F9** to check basic syntax. If one suspects a schema-validity problem but can't quite track it down by eye, one can always add the `schemaLocation` phrases for hand-editing the files, remembering to comment them out before attempting deployment of the web application.



Ant Build Process

Though most students will be happy to let the **ant** command take care of things, some students (and most instructors) will want to understand the inner workings here a little better. Each web project in the course has its own **build.xml** and **build.properties** files; these rely on the master environment variable **CC_HOME** to import targets and properties in a central directory, **%CC_HOME%\Spring\Ant**. Information here defines a routine for building a Spring application; we won't document this in detail here, but rather summarize the interesting tasks:

- Clean out and create a **build** directory, which becomes the root of a web application (cleanup is a little tricky since Tomcat likes to lock certain files)
- Compile Java source files to **build\WEB-INF\classes**
- Copy resources from **docroot** to **build** – this may include HTML and JSP, the **web.xml** and Spring configuration files, message resources, and custom tags
- Copy the Spring JAR into **build\WEB-INF\lib**
- Create a Tomcat context file, e.g. **LandUse.xml** for the LandUse application, and copy it into Tomcat's **conf\Catalina\localhost** directory

Not everyone is familiar with this last step, or with this approach to deploying applications to Tomcat. Note that there is no WAR file, nor do we truly deploy our code to Tomcat. Rather, we use the context file to tell Tomcat where to find our application, and to map it to a certain context-root URL. This has the advantage of keeping our build entirely local and letting Tomcat reload as necessary when we change a JSP or class file. It can lead to some confusion, though: once the context XML is in Tomcat's tree, a build failure will leave Tomcat thinking that the application is deployed when it's not even well-formed. This may manifest as exceptions showing in Tomcat's console on a later startup. It's typically harmless, and easy to fix: just rebuild the application completely, from any of its example, demo, or lab steps; the build will overwrite the current context XML to point to itself, and Tomcat will start serving files from that recent build.





Eclipse Overlays

Capstone Courseware provides an optional package of workspace and project files for Eclipse WTP 1.5 for this course. (See the course Setup Guide for download URLs.)

Instructors, use this package on your own initiative and at your own risk. You should have significant experience yourself with Eclipse WTP before using the overlay package in the classroom. The workspace and projects have been tested lightly with the course but are not part of the standard product.

That said, this overlay should save a good deal of work for those who wish to use Eclipse instead of the text editor and command-line tools that are standard for the course. See the documentation inside the workspace itself for general notes on usage: when you first open the workspace you'll see a **ReadMe.txt** that directs you to deeper, HTML-based documentation. Be prepared to walk students through the first few exercises in Eclipse; the notes in this file are for experienced Eclipse users, and will not be clear to many students on their own.

See also the file **Spring Module Notes.html** (available in the IDE as well) for specific matters related to using Eclipse WTP with this course.

Again, Capstone Courseware can only offer complete technical support on the standard course, and while we hope this overlay is convenient, it is not as thoroughly tested as the core lab image at this time. If a given exercise is giving trouble, please be certain to build and run it from the command line, using the SDK tools as prescribed in the student guide, before contacting Capstone.





Teaching Notes

Chapter 1

This chapter should accomplish two things: give students a straightforward overview of Spring, and prove the machine setups by building one each of the two primary types of applications in the course, which are a standalone application that uses a bean factory and a web application that uses a web context configuration. Along the way students can get familiar with the layout of course exercises. Don't let inquisitive students bog you and the rest of the class down during these two examples; there's plenty of time to delve into all the details in the next few chapters, and you'll be seeing both of these applications again.

Chapter 2

Here we start building the foundation of Spring IoC, which will gradually become central to course exercises over the remaining chapters, as it is central to most Spring development. At first the beauty of the bean factory, inversion of control, and dependency injection may not be obvious to everyone. By the end of the class it probably will, but in the meantime, try to promote the idea that while J2EE gives a set of IoC features to a select few types (servlets, EJBs, etc.), Spring is egalitarian: it gives these same features to any Java class. How nice to be able to configure any class you write with properties that might not be known at development time! It's the end of the properties file, ad-hoc system properties, and so forth, and much more elegant and powerful than any of those hoary mechanisms.

An interesting side note that you might make during or after the Singleton vs. Singleton exercise: in order to enforce singularity (or prototype behavior) even over the rules set by a Java developer, Spring has to carry out some privileged actions. For this to work, either there must be no security manager in place, or it must have a policy that allows the Spring codebase some special permissions. A rough list is here:

```
permission java.io.FilePermission "", "read";
permission java.io.FilePermission "-", "read,write";
permission java.util.PropertyPermission "*", "read";
permission java.net.SocketPermission
    "www.springframework.org:80", "connect";
permission java.lang.RuntimePermission "accessDeclaredMembers";
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";
```



Chapter 3

Though technically it introduces the distinct idea of dependency injection, this chapter is largely an extension of Chapter 2. We carry our study of Spring IoC to completion, and by the end we should be able to construct arbitrarily complex graphs of objects: for the business tier, the web tier, or the persistence tier, later in the course. Remember the Spring IoC mantra of using the lightweight container to “create, configure, and assemble.” Chapter 2 is about “create,” and a little about “configure;” this one’s more on “configure” and a lot about “assemble.”

One note on the Policy demo: to accommodate the different concrete paths at which the code is deployed, later examples will rely on the `CC_HOME` environment variable, and various beans will on their own initiative expand the phrase “`{env.CC_HOME}`” found in their property values by looking up this variable. But that seems a bit much for this exercise, and so the exact path of the code source must be entered when configuring the policy. This is simple enough and the path is shown in the demo instructions. But the demo answer will have a path from the **Examples** tree, so any copy-and-paste solution will need some additional editing. It also means that course deployments to non-standard roots (other than `c:\Capstone`) will require editing of the XML file, even for the answer code.

Chapter 4

This chapter is the big brick in the course. There’s a lot to cover that all seems equally fundamental, and so here it all is: the dispatcher servlet, controllers, handler mappings, views, resolvers, and configuration. Don’t worry if this one seems to go on quite a while; it will be time well spent, and rushing this material certainly won’t serve students well, since the next few chapters build on what’s here. Do plan to chop this one up into sections, with a break or two in between. Then give plenty of time for the labs at the end.

You may well have your own (even contradictory) advice to offer at the end of the chapter, in the “Which Way to Go?” discussion – so much the better. This is entirely subjective stuff, and if there were one best way to map URLs to controllers or to manage the dependencies between controllers and views, then presumably Spring would stick with that and leave out the others.





Chapter 5

This chapter and Chapter 6 are tightly linked, and it was tricky to figure out exactly how to divide and organize these topics. Basically we're trying to establish command objects first, and the full flow orchestration that comes with form controllers in the next chapter. Then, the **MultiActionController** needs a home, and this seems like the right place, especially as it gives some exercise in dealing with command objects and validators more or less manually.

MultiActionController is a bit of an odd duck, actually, living off to the side of the main line of inheritance in the servlet-MVC package. This and the next chapter are mostly a long walk down that inheritance tree, and that means this class is a sort of side road.

Chapter 6

In this chapter, we try to tackle just one main job, which is managing HTML forms with form controllers – but it's a tricky task, to say the least. The understatement of the whole coursebook may be that **SimpleFormController** isn't really simple. On the one hand this class is the crown jewel of the MVC library, and has great facility once you get control of it; on the other hand it's a remarkably hard beast to tame. The demos are intended to walk students through the many necessary steps so as to make the task not too daunting.

Otherwise, this is one of those tools with which it's best to work from a successful example and start varying things, rather than build from scratch and watch things fail and fail until suddenly they work. Be sure you know these demos very well before tackling them in class.

By the way, this may be the right place to point out a custom tag that's been in use in the course examples since Chapter 4. There's enough going on in that chapter that it seems best to leave `<cc:springBean>` undiscovered for a little while; but if students ask, or once you reach this point in the course, you might want to dig into it and explain what it's doing and how it's doing it. The idea here is to bring Spring beans into a JSP – even if that JSP doesn't enjoy the services of the Spring MVC framework. So for the database wrapper in LandUse, and similarly in the Wholesale application, the JSPs use this tag to get access to an object that all the other components in the application load using the application context.





Chapter 7

This chapter closes out coverage of the web module, and it may be seen as a bit of a hodgepodge of advanced topics. But there is a clear theme, or at least there are a few! Exception handlers and interceptors more or less of a piece, and the context-and-lifecycle section seems to put the finishing touches on the idea that any JavaBean can have the first-class treatment of a contained object. But ultimately this is the stepping-off point, and towards the end of the chapter we're mostly telling students where they might look next as they continue to learn about Spring.

Chapter 8

Responding to the demand as we see it at this time, this course goes heavy on the web module and very light on the persistence support in Spring. There is certainly plenty more to study on this end of a Java EE application, and a five-day course that gives the web and the persistence tiers each a couple of days may someday emerge. But for one thing, most audiences seem to mean "Spring MVC" when they say "Spring;" and for another, it's hard to go much deeper into Spring's persistence framework without studying Hibernate or one of its competitors. This course is designed in part to accompany and lead into our two-day course on Hibernate, so we keep things fairly sketchy in this chapter.

The highest value pieces of this chapter are the JDBC template and the declarative transactions support. We do a lab in one and finish with a demo in the other; we could get deeper into transactions, but more and more it seems that any lab that concludes a course is doomed never to be done by students; so we're trying here for a mid-chapter lab that will get done, and a demo that makes a nice closing act for the instructor.

One note about the database files: for simplicity of setup and use, we've deployed Apache Derby as our RDBMS for this chapter. Derby's nice, and simple to work with, but in embedded mode it suffers a severe limitation, which is that only one JVM can connect to a schema at a time. This doesn't pose a problem in these exercises until the web application is deployed with the JDBC DAO; from that point on, if you want to, say, run the standalone application that refreshes the data in the table, you have to shut down Tomcat to do it, then start Tomcat up again to test the web application. This is usually no more than a minor irritation in practice, but it's good to know about it in advance; otherwise the exceptions that you see will make it seem like something really nasty is going on and students may be tempted to bail on the demo. In a later revision of the course, we may move to another RDBMS or use Derby's network driver and a separate RDBMS process.



Revision History

Version 2.0.1 fixes a handful of typos and omissions in the book. There are no changes to the lab code, but the Eclipse overlay is updated from the old style to the new one that supports WTP 1.5.

Version 2.0 is the initial release of the course.





Troubleshooting

If you run into any trouble with code exercises, the first and best thing to do is to double-check that the classroom machines have been set up precisely according to the course setup guide. Especially, the wrong version of a tool can cause significant problems; don't wander off-book in this way unless absolutely sure you can support the software that you prefer and that we haven't tested. Check environment variable settings carefully, too; these are the cause of a great many classroom glitches.

Below are some specific pitfalls that have come up in previous offerings of the course:

- If you set the **CC_HOME** environment variable in an individual command console (and not globally, e.g. through the Windows Control Panel), you must be sure to start Tomcat from that console or from one that also has the environment variable defined correctly. This is necessary because some Java classes running in the deployed web applications will attempt to resolve the value of that variable, as a way of finding resources outside of their own web contexts. A good strategy is to set your environment in the console where you'll do your Ant builds; **start** another console, which inherits that environment, and change to the Tomcat **bin** directory in that second console. Use the second console to **startup** and **shutdown** Tomcat as necessary.
- Attempts to validate Spring beans XML files against the schema will succeed only if there's an Internet connection, since the **schemaLocation** attribute points to the public schema. If the classroom is not networked, you can reset this attribute to the schema cached locally in **%CC_HOME%\Tools\Spring2.0\dist\resources**. (This change would have to be made for each XML file, which is a bit clunky, but once you've done it, it's easy to do again by copy-and-paste.)





Errata

The following issues have arisen since the release of this revision:

- On page 11, the directory listing shows **Checking.java**, which is actually **CheckingAccount.java**.





Feedback

We very much appreciate whatever feedback we can get on our courseware – especially from the instructor's perspective. Naturally, the more specific, the better, and we strongly encourage you to make notes on issues you may encounter in the classroom, whether they're typos, missing files, or suggestions for clearer language to explain a concept. We can't guarantee that we'll act on every suggestion, but we're aggressive about stamping out problems and try to be highly responsive. Hopefully this means that when you give us good feedback, you get a better course the next time you need to teach it.

Please direct all courseware feedback to

Will Provost
Capstone Courseware
<mailto:provost@capcourse.com>
877-227-2477

For anyone who's interested, we have a very informal defect-tracking system, based in Excel spreadsheets with columns to capture defect location, nature, status, and author feedback. Ultimately, feedback goes into these sheets, so if you want a template, we'll be happy to provide one, to facilitate the reporting process.

