

JavaTM magazine

By and for the Java community 



CDI 34

JYTHON 42

DOCKER PRIMER 51

WEBSOCKETS 58

NOVEMBER/DECEMBER 2015

Libraries

FINDING THE RIGHT ONE

13 JCOMMANDER:
COMMAND-LINE
PARSING

19 BYTE BUDDY:
BYTECODE
GENERATION

24 JSOUP: ELEGANT
HTML PARSING

30 HOW THE JVM
FINDS AND LOADS
LIBRARIES

An easy-to-use library for parsing
the most-complex command lines

COVER ART BY I-HUA CHEN

XRebel

 ZEROTURNAROUND

THE LIGHTWEIGHT JAVA
PROFILER

TRY IT FREE NOW!

*Get a free
t-shirt! →*



However, the pool of widely used OSI-approved licenses bulges with numerous provisions, many of which require an attorney to understand. Several OSS license analysts keep databases

But without wasting time on websites, in forums, and with

In software, this needs to be implemented—just a simple, sane system devoid of complex, competing, historical artifacts.

CREATE
THE FUTURE

oracle.com/java

20
YEARS

1995-2015

Java™

ORACLE®

Why Compile Java Applications to Native Code Executables?

Since the early years of Java's two-decade history, production-grade JVMs relied on just-in-time (JIT) compilation for performance. A JIT compiler kicks in at application runtime and compiles the "hot" methods detected by the bytecode interpreter to native code.

The problem is that upon termination the JVM discards all code produced by the JIT compiler, which means it has to start over next time. The overhead of the interpret-profile-compile

developer's system, although Android ART performs native compilation directly on the target device at application install time.

AOT compilation eliminates the warm-up cycle, enabling a large Java application to start 2 to 3 times faster and run at full speed from the start. It can also deliver an overall performance boost, especially in constrained environments with no spare computational resources or battery power for an advanced JIT compiler, such as embedded devices.

However, while faster startup and lower latency are valuable benefits, neither is the main motivation for using an AOT compiler when developing Java applications for desktop and server platforms.

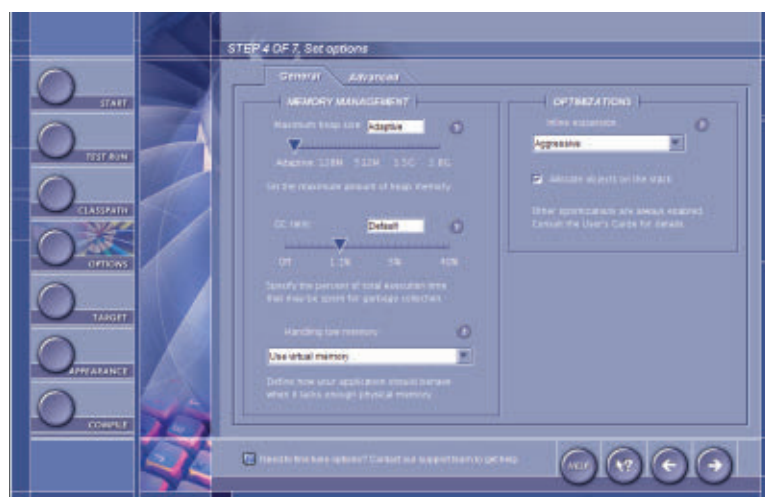
The unique advantage of compiling the *entire* Java application to native code before deploying it in production or shipping to end users is actually *protection against Java decompilers*.

Class files emitted by the standard javac compiler are extremely easy to revert back to source—just search for "download java decompiler" and you will get your source code back in five minutes. Methods exist for making decompiler output less comprehensible, but they have serious limitations and may negatively affect application robustness and performance. Name obfuscation can only be applied to application classes that are not accessed via reflection or JNI, whereas references to the standard library classes remain visible. Excessive control flow

obfuscation substantially hinders JIT compiler optimizations, penalizing application performance. Finally, software encryption does not protect your Java classes at all—they must be decrypted prior to execution, and dumping classes to disk as the JVM loads them is fairly trivial.

In contrast, a native binary produced by an optimizing AOT compiler carries no performance tax and is about as difficult to reverse-engineer as if you had coded the original program in C++. And you can still take additional measures to protect your intellectual property or enhance the security of your application if you wish: obfuscate names, encrypt strings and resources, and further protect the native code executable using platform-specific tools.

Excelsior JET 11 is the only certified Java SE 8 implementation capable of precompiling Java applications to optimized native code executables. It currently supports Windows, OS X, and Linux on Intel hardware.



Excelsior JET Control Panel

process, known as the "warm-up cycle," makes Java applications start more slowly, with slower initial response times compared to functionally equivalent native applications. This leads to the idea of ahead-of-time (AOT) compilation.

An AOT compiler precompiles some or all application classes to native code before the application is run. That usually happens on the

For more information, visit www.ExcelsiorJET.com



JULY/AUGUST 2015

I enjoyed the July/August issue. However, I disagree with the answer to question 2 in Fix This. The article said the answer is B, but the answer is A. When the code executes, a single object is created on line 4, and at the end of the execution of the `main` method it is eligible for garbage collection because the created object is only ever referenced by local variables `e` and `e1`.

In Fix This for the July/August issue, question 3 contains an error. It refers to a code snippet in a file called Shop.java, which declares a class entitled `OnlineCart`. This is not legal in Java.

Editor Andrew Binstock responds: As mentioned in the introduction to this quiz, the questions were provided by the team that supports the Java certification exam publishing. After discussion with them, it came to light that the quiz questions had not been through the required technical review. The process has now been modified to ensure that only thoroughly reviewed quiz questions are

I read your review of Paul and Gail Anderson's book about JavaFX and the NetBeans Platform in the September/October issue. In commending the publishers for printing a book for "so narrow a topic," you ask, "How many such readers could there be?"

I noticed in the September/
October issue that one of your

Concerning Jennifer Hamilton's reply in the September/October issue to Raj Thondepu, who inquired about a paper edition of *Java Magazine*: has Oracle actually researched this before concluding that a printed magazine wouldn't be sustainable?

We welcome comments, grumbles, and kudos at javamag_us@oracle.com. These might be edited for publication. If your note is private, please so indicate. For other ways to reach us, see the last page of this issue.



Devoxx Belgium NOVEMBER 9–13

ANTWERP, BELGIUM

By developers for developers, this event has 200 speakers and 3,500 attendees from 40 countries. Tracks this year include Java SE, JVM languages, and server-side Java, as well as cloud and big data, mobile, and architecture and security, among others.

W-JAX 15

NOVEMBER 2–6

MUNICH, GERMANY

The W-JAX conference covers current and future-oriented technologies from Java, Scala, Android, and web technologies to agile development models and DevOps.

J-Fall 2015

NOVEMBER 5

EDE, NETHERLANDS

The annual Java conference organized by the Dutch Java User Group (NLJUG) typically sells out and has outgrown its usual venue. This year, J-Fall will take place in the CineMec in Ede.

Devoxx Morocco

NOVEMBER 16–18

CASABLANCA, MOROCCO

Formerly the JMaghreb conference, this event is a university day of training, workshops, and labs followed by conference days of sessions on software development, web, mobile, gaming, security, methodology, Internet of Things, and cloud. The Decision Makers evening includes discussion of issues related to the IT industry in Morocco.

QCon San Francisco 2015

NOVEMBER 16–20

SAN FRANCISCO, CALIFORNIA

A practitioner-driven software development conference, QCon is designed for technical team leads, architects, engineering directors, and project managers who influence innovation in their teams. Tracks this year include Taking Java to the Next

Level and The Dark Side of Security. The last two days are devoted to workshops.

Codemotion Milan

NOVEMBER 18–21

MILAN, ITALY

This conference is open to users of all languages and platforms. It offers full-day workshops on the first two days, followed by keynotes and conference sessions.

Codemotion Spain

NOVEMBER 27–28

MADRID, SPAIN

This two-day event draws 2,000 attendees, represents more than 30 communities, and features coding lectures and workshops. Activities for startups, recruiting, and networking are included.

Clojure eXchange 2015

DECEMBER 3–4

LONDON, ENGLAND

Meet with the world's leading experts, learn how to use Clojure with your team, and discuss war stories with your peers. Both days will feature a mixture of talks covering various aspects of Clojure devel-



opment: from libraries to music, from ClojureScript to data.

Groovy and Grails eXchange 2015

DECEMBER 14–15

LONDON, ENGLAND

Stay ahead of the curve and hear the 2016 roadmap for Groovy and Grails from core commit-
ters and Groovy authorities
Guillaume Laforge and Graeme
Rocher. Engage with other lead-
ing experts and fellow enthusiasts
and learn the latest innovations
and practices.

Jfokus

FEBRUARY 8–10, 2016

STOCKHOLM, SWEDEN

Jfokus has run for eight years

and is the largest annual Java
developer conference in Sweden.
Conference topics include Java SE
and Java EE, front end and web,
mobile, continuous delivery and
DevOps, Internet of Things, cloud
and big data, future and trends,
alternative JVM languages, and
agile development.

DevNexus 2016

FEBRUARY 15–17, 2016

ATLANTA, GEORGIA

DevNexus is a conference draw-
ing 1,700 developers, with 6 work-
shops, 12 tracks, and 120 presen-
tations. Featured tracks include
HTML5 and JavaScript, Java SE/
Java EE/Spring, and data and
integration.

ConFoo 2016

FEBRUARY 22–26, 2016

MONTREAL, QUEBEC, CANADA

ConFoo is a multitechnology con-
ference for web developers, fea-
turing about 150 presentations
by popular international speak-
ers. Past sessions have included
Testing Java EE Applications Using
Arquillian by Reza Rahman and
Hybrid Mobile Development with
Apache Cordova and Java EE 7 by
Ryan Cuprak.

Embedded World 2016

FEBRUARY 23–25, 2016

NUREMBERG, GERMANY

The 14th annual gathering of
embedded system developers
will explore the latest develop-
ments, define trends, and once
again present the key areas of
focus for future developments.
This is where hardware, software,
and system development engi-
neers come together to turn the
next milestones of the Internet of
Things into reality.

Apache Hadoop Innovation Summit

FEBRUARY 25–26, 2016

SAN DIEGO, CALIFORNIA

With presentations from more
than 25 hands-on industry speak-

ers, topics covered will include
MapReduce and Spark, building
privacy-protected data systems,
scalable data curation, best prac-
tices, and architectural consider-
ations for Hadoop applications.

Riga Dev Day

MARCH 2–4, 2016

RIGA, LATVIA

This event is a joint project by
Google Developer Group Riga, Java
User Group Latvia, and Oracle User
Group Latvia. By and for software
developers, Riga Dev Day focuses
on 25 of the most-relevant topics
and technologies for that audi-
ence. Tracks include JVM and web
development, databases, DevOps,
and case studies.

Have an upcoming confer-
ence you'd like to add to our
listing? Send us a link and a
description of your event at
least four months in advance at
javamag_us@oracle.com. We'll
include as many as space permits.



By Stephen Chin and James Weaver
Oracle Press


Most projects built today on the Raspberry Pi are undertaken by hobbyists, in part because of its remarkable ease of use: only a modicum of knowledge is required to get up and running. The hardware is comparatively straightforward, and the key obstacles consist principally of understanding the interactions between the software and the hardware. This book is an intro-

There follow projects that educate the reader on the basic use

My only complaint is that this book lacks an appendix containing reference information, so if you want to refer back to some hardware details, you need to remember in which project it was first presented. Other than this minor detail, this book is by far the best introduction I've seen to Java programming on the Raspberry Pi. —*Andrew Binstock*

Your Destination for Java Expertise


Written by leading Java experts, Oracle Press books offer the most definitive, complete, and up-to-date coverage of Java available.



Raspberry Pi with Java: Programming the Internet of Things (IoT)

Stephen Chin, James Weaver


Use Raspberry Pi with Java to create innovative devices that power the internet of things.



Introducing JavaFX 8 Programming

Herbert Schildt


Learn how to develop dynamic JavaFX GUI applications quickly and easily.



Java: The Complete Reference, Ninth Edition

Herbert Schildt

Fully updated for Java SE 8, this definitive guide explains how to develop, compile, debug, and run Java programs.



OCA Java SE 8 Programmer I Study Guide (Exam 1Z0-808)

Edward Finegan, Robert Liguori

Get complete coverage of all objectives for Exam 1Z0-808. Electronic practice exam questions are included.

Syntax Flexibility

To support as many syntax styles as possible, JCommander allows you to specify separators other than spaces. For example, instead of `java Main -log 3` you might want to use `java Main -log=3` or `java Main -log:3`, which can all be specified with the `separators` attribute:

```
@Parameters(separators = "=")
public class SeparatorEqual {
    @Parameter(names = "-level")
    private Integer level = 2;
}
```

Validation

As your command grows complex, it will become tricky to decide whether a given command line is valid, because multiple options can interact with each other in various ways. JCommander can provide some assistance here by making it easy to validate your parameters. The syntax is very similar to the one we just covered with type converters:

```
@Parameter(names = "-age",
            validateWith = PositiveInteger.class)
private Integer age;
```

And here is the implementation of this validator:

```
public class PositiveInteger
    implements IParameterValidator {
    public void validate(String name, String value)
        throws ParameterException {
        int n = Integer.parseInt(value);
        if (n < 0) {
            throw new ParameterException(
                "Parameter " + name + " should be positive"
                + (found " + value +"));
        }
    }
}
```

}

Complex Commands

You might be familiar with a few tools that use subcommands to express sophisticated invocation syntax. For example, `git` offers this kind of syntax when multiple subcommands have their own syntax: if you call `git commit`, then you can specify parameters such as `--author` or `-amend`, while `git add` accepts `-i`. It's easy to implement this kind of syntax with `JCommander`.

First of all, you define your commands in their own classes. Here is `commit`:

```
@Parameters(separators = "=",
            commandDescription = "Record changes")
private class CommandCommit {

    @Parameter(description = "The list of files")
    private List<String> files;

    @Parameter(names = "--amend", description =
        "Amend")
    private Boolean amend = false;

    @Parameter(names = "--author")
    private String author;
}
```

And here is `add`:

```
@Parameters(commandDescription =
    "Add file to the index")
public class CommandAdd {

    @Parameter(description =
        "File patterns for the index")
    private List<String> patterns;

    @Parameter(names = "-i")
```



```
private Boolean interactive = false;
}
```

Then, you add these commands to JCommander. In the following code, I have added a few assertions at the end to demonstrate what happens:

```

JCommander jc = new JCommander();

CommandAdd add = new CommandAdd();
jc.addCommand("add", add);
CommandCommit commit = new CommandCommit();
jc.addCommand("commit", commit);

jc.parse("-v", "commit", "--amend",
        "--author=cbeust", "A.java", "B.java");

Assert.assertTrue(cm.verbose);
Assert.assertEquals(jc.getParsedCommand(),
        "commit");
Assert.assertTrue(commit.amend);
Assert.assertEquals(commit.author, "cbeust");
Assert.assertEquals(commit.files,
        Arrays.asList("A.java", "B.java"));

```

As you can see from the assertions, the code has parsed the command line correctly and placed the arguments in the expected variables.

Architecture

JCommander can support the most-complex codebases and syntax styles, so there are several other features that help you organize your code cleanly.

Multiple argument objects. As your syntax grows, you might find yourself having one gigantic argument class that becomes a bit difficult to maintain. JCommander lets you break this class into multiple classes so that you can organize the options in a more intuitive way:

```
CommandRead argRead = new CommandRead();
CommandWrite argWrite = new CommandWrite()
JCommander jc = new JCommander(argRead, argWrite);
jc.parse(argv);
```

```
// argRead and argWrite are now both initialized
```

Parameter delegates. As you write multiple programs, you might find yourself wanting to reuse existing arg classes, which is something you can do with parameter delegates. In short, parameter delegates are pointers to other arg classes. In the previous example, I decided to create two different arg classes and declare these directly in JCommander; but instead, I might want to delegate to them. This is done with the `@ParameterDelegate` annotation:

```
class MainParams {
    @Parameter(names = "-v")
    private boolean verbose;

    @ParametersDelegate
    private ArgRead argRead = new ArgRead();

    @ParametersDelegate
    private ArgWrite argWrite = new ArgWrite();
}
```

After this declaration, I need to declare one argument parameter, `MainParams`, and it will contain the aggregation of both `ArgRead` and `ArgWrite`.

Polyglotism

Thanks to the JVM's ability to support multiple languages, JCommander is trivial to use from any JVM language. I'm currently using it on a Kotlin project:

```
class Args {
    @Parameter(names = arrayOf("--buildFile"))
```


sented by `MethodDelegation`: When delegating to a method, you can first execute your custom code, and then delegate the call to the original implementation. Additionally, you can also dynamically access the information of the original call site using the `@Origin` annotation, as shown in Listing 2. When delegating to other methods, you can also dynamically access the information of the original call site, as shown next.

■ **Listing 2.**

```
public static class Agent {
    public static String record(@Origin Method m) {
        System.out.println(m + " called");
    }
}
```

```
Class<?> clazz = new ByteBuddy()  
    .subclass(Object.class)  
    .method(ElementMatchers.isConstructor())  
    .intercept(MethodDelegation  
        .to(Agent.class)  
        .andThen(SuperMethodCall.INSTANCE))  
    // & make instance;
```

`MethodDelegation` automatically looks up the best match of method signatures in case multiple interception targets are available. While the lookup is powerful and can be customized, I recommend keeping the lookup simple and understandable. After the method has been invoked, the original call continues, thanks to `andThen(SuperMethodCall.INSTANCE)`.

The target method can take a couple of annotated parameters. To access the arguments of the originating method, you can use `@Argument(position)` or `@AllParameters`. To obtain information about the originating method itself, you can use `@Origin`. The type of that parameter can be `java.lang.reflect.Method`, `java.lang.Class`, or even `java.lang.invoke.MethodHandle` (the latter, if used with

Java 7 or later). These arguments provide information about where the method has been called from, which could be useful for debugging, or even about taking different code paths, in the event that the same method is an interception target for multiple methods.

To call the originating method or its super method from the target method, Byte Buddy provides `@DefaultCall` and `@SuperCall` parameters.

Mocking

Sometimes you want to write a unit test for a scenario that can happen at runtime, but you cannot provoke that scenario reliably for the purpose of the test, if at all. For instance, in **Listing 3**, the random number generator needs to produce a specific result for you to test the control flow.

■ Listing 3.

```
public class Lottery {
    public boolean win() {
        return random.nextInt(100) == 0;
    }
}
```

```
Random mockRandom = new ByteBuddy()  
    .subclass(Random.class)  
    .method(named("nextInt"))  
    .intercept(value(0))  
    // & make instance;
```

```
Lottery lottery = new Lottery(mockRandom);
assertTrue(lottery.win());
```

Byte Buddy provides various kinds of interceptors, so writing mocks, or spies, is easy. However, for more than a few mocks, I would recommend switching to a dedicated mocking library. In fact, version 2 of the popular mocking library [Mockito](#) is currently being rewritten to be based on Byte Buddy.

classes it should `rebase`. This example will modify only classes having the `javax.ws.rs.Path` annotation. Next, I tell the builder how to `transform` those classes. In this example, the agent will intercept calls to either `GET` or `POST` annotated methods and delegate to the `profile` method. For this to work, the agent needs to be hooked into the Instrumentation using `installOn()`.

The profile method itself uses three annotations: `RuntimeType`, to tell Byte Buddy that the return type `Object` needs to be adjusted to the real return type used by the method it intercepts; `Origin`, to obtain a reference to the actual method intercepted, which is used to print its name; and `SuperCall`, to actually perform the original method call. In contrast to the previous example, I need to perform the super call myself, because I want to be able to have my code executed before and after the method call—so that I can perform the timing.

Comparing the way Byte Buddy implements method interception to the default Java `InvocationHandler`, you can see that the Byte Buddy method is much more optimized due to the fact that the interception will pass in only the required arguments, while `InvocationHandler` must fulfill the following interface:

```
Object invoke(Object proxy,
              Method method, Object[] args)
```

This benefit is especially noticeable for primitive arguments or return types, which need to be autoboxed. The additional `RuntimeType` annotation causes Byte Buddy to reduce any boxing to a minimum. Even though the JVM mostly optimizes away simple boxings, this is not always true for complex interfaces such as that of the `InvocationHandler`.

Using an Agent Without -javaagent

Using an agent to generate and modify code at runtime is a

powerful technique; however, forcing the `-javaagent` argument to make it work is sometimes inconvenient. Byte Buddy comes with a handy convenience feature that uses the [Java Attach API](#), which originally was designed to load diagnostic tooling at runtime. It attaches the agent to the currently running JVM. You need the additional `byte-buddy-agent.jar` file, which contains the utility class `ByteBuddyAgent`. With that, you invoke `ByteBuddyAgent.installOnOpenJDK()`, which does the same thing that starting the JVM with `-javaagent` did. The only other difference with this approach is that you do not invoke `installOn(inst)`, but rather you invoke `installOnByteBuddyAgent()`.

Conclusion

Despite the existence of dynamic proxies in the JDK and three popular, third-party, bytecode-manipulation libraries, Byte Buddy fills an important gap. Its fluent API uses generics, so you do not lose track of the actual type you are modifying, which can easily happen using other approaches. Byte Buddy also comes with a rich set of matchers, transformers, and implementations, and it enables their use via lambdas, which results in relatively concise and readable code.

As a result, Byte Buddy is fully understandable by developers who are not accustomed to reading bytecodes and working at low levels. With the upcoming version 0.7, Byte Buddy will support all the infrastructure around generic types. This way, Byte Buddy allows for easy interaction with generic types and type annotations even at runtime. As someone who writes a lot of bytecode-handling code, I both recommend and use this library. [Byte Buddy received a Duke's Choice Award at this year's JavaOne conference. —Ed.] [</article>](#)

LEARN MORE

- [JVM Specification for Java 8](#)
- [Byte Buddy on Stack Overflow](#)



Easily parse HTML, extract specified elements, validate structure, and sanitize content.

What It Is

jsoup can parse HTML files, input streams, URLs, or even strings. It eases data extraction from HTML by offering document object model (DOM) traversal methods and CSS and jQuery-like selectors.

It can manipulate the content: the HTML element itself, its attributes, or its text. It also updates older content based on HTML 4.x to HTML5 or XHTML by converting deprecated tags to new versions. It can also do cleanup based on white-lists, tidy HTML output, and complete unbalanced tags

automagically. I will demonstrate these features with some working examples shortly.

All the examples in this article are based on jsoup version 1.8.3, which is the latest available version at the time of this writing. The complete source code for the article is available on GitHub.

The DOM and jsoup Essentials

DOM is the language-independent representation of the HTML documents, which defines the structure and the styling of the document. **Figure 1** shows the class diagram of jsoup framework classes. Later, I'll show you

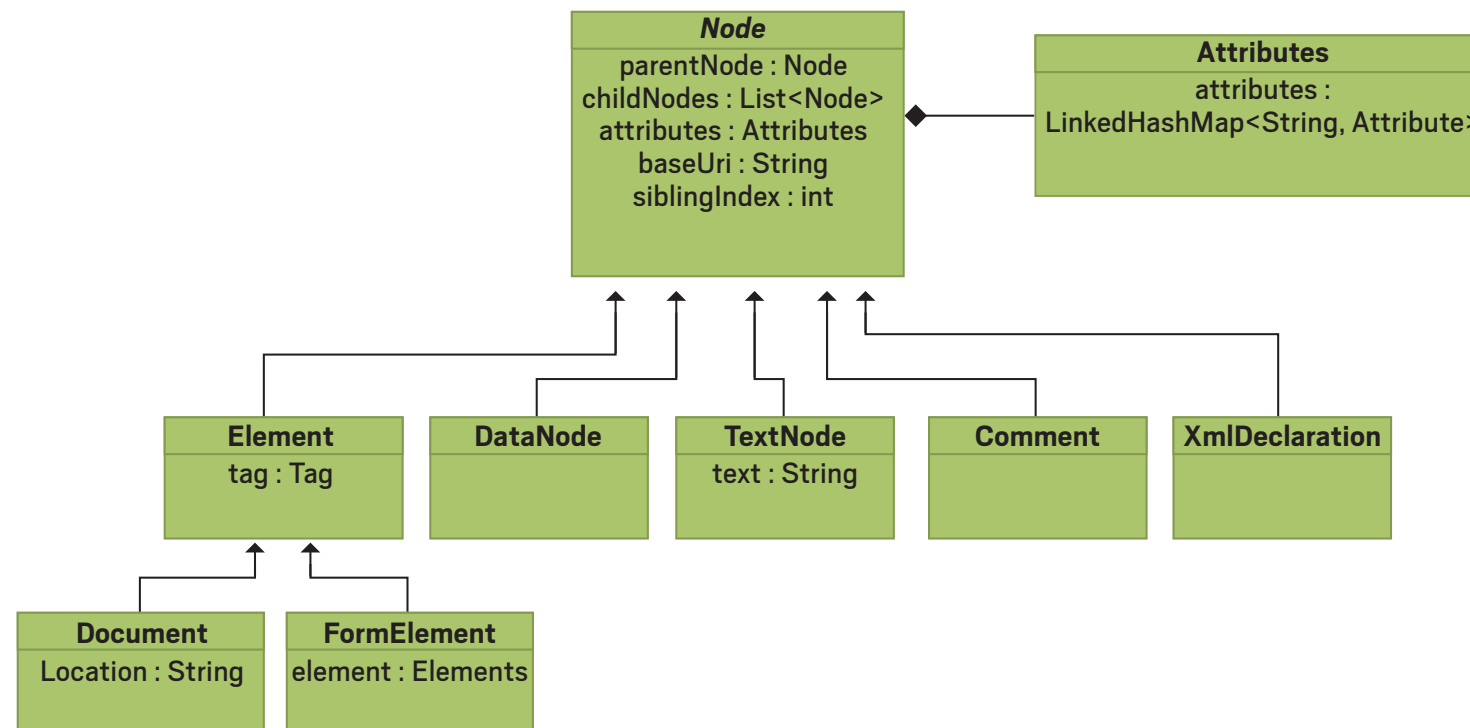


Figure 1: jsoup class diagram

and jQuery-like selectors. I will demonstrate both approaches by parsing a web page and extracting all links that have HTML `<a>` tags. The code in **Listing 2** parses the Java Champions bio page and extracts the link names for all the Java Champions marked as “New!” (see **Figure 2**).

The marking was done by adding a `` tag with text **New!** right next to the link. So, I will be checking for the content of the next-sibling element of each link.

■ **Listing 2.**

```
public class Example2Main {

    public static void main(String... args)
        throws IOException {
        Document document = Jsoup.connect(
            "https://java.net/website/" +
            "java-champions/bios.html" )
            .timeout(0).get();

        Elements allElements =
            document.getElementsByTag("a");
        for (Element element : allElements) {
            if ("New!".equals(
                element.nextElementSibling()!=null
                ? element.nextElementSibling()
                .ownText()
                : "")) {
                System.out.println(
                    element.ownText());
            }
        }
    }
}
```



Figure 2: Part of the HTML page to be parsed

}

The same extraction of the links can also be done with selectors, as shown in **Listing 3**. It extracts the links that start with href value #.

■ **Listing 3.**

```
public class Example3Main {  
  
    public static void main(String... args)  
        throws IOException {  
        Document document = Jsoup.connect  
            ("https://java.net" +  
             " /website/java-champions/bios.html")  
                .timeout(0).get();  
        Elements allElements = document.select  
            ("a[href*=#]");  
        for (Element element : allElements) {  
            if ("New!".equals(element  
                                .nextElementSibling() != null  
                                ? element.nextElementSibling  
                                  ().ownText() : "")) {  
                System.out.println(element  
                                    .ownText());  
            }  
        }  
    }  
}
```

Selectors are powerful compared with DOM-specific methods. They can be combined together to refine selection. In the previous code examples, we are doing the “New!” text check by ourselves, which is trivial. The example in **Listing 4** selects the `` tag that contains the “New!” text, which resides after a link that has an `href` starting with the value `#`. This really shows the power of selectors.

jsoup also offers useful features such as tidying HTML and blending the structure or pseudo-structure of CSS selectors.

allows these HTML tags: a, b, blockquote, br, cite, code, dd, dl, dt, em, i, li, ol, p, pre, q, small, span, strike, strong, sub, sup, u, ul.

```
@Test
public void basicCleaningWorksOK() {
    String html = "<div><p><a " +
        "href='javascript:hackSystem()' " +
        "'>Hello</a></div>";
    String cleanHtml = Jsoup.clean(html,
        Whitelist.basic());
    assertThat(cleanHtml, is("<p><a " +
        "rel=\"nofollow\">Hello</a></p>"));
}
```

Conclusion

In this article, I have shown only a subset of what jsoup can do. It also offers useful features such as tidying HTML, manipulating HTML tags' attributes or texts, and blending the structure or pseudo-structure of CSS selectors. Put another way, any HTML processing you might need to do is a likely candidate for using jsoup. </article>

```
@Test
public void simpleTextCleaningWorksOK() {
    String html = "<div>" +
        "<a href='http://www.oracle.com'>" +
        "<b>Hello + Reader</b>!</a></div>";
    String cleanHtml = Jsoup.clean(
        html, Whitelist.simpleText());
    assertThat(cleanHtml,
        is("<b>Hello Reader</b>!"));
}
```

Listing 9 shows an example of the usage of `basic()`, which

- [jsoup on Stack Overflow](#)
- [jsoup recipes](#)
- Comparison of HTML parsers

of the `java.lang.ClassLoader` type that developers can use to locate and load the classes by name. If you're confused by this chicken-and-egg problem and wonder how the first class loader that loads all the JDK classes (for example, `java.lang.String`) is created, you're thinking along the right lines. Indeed, the primordial class loader, called the *bootstrap class loader*, comes from the core of the JVM and is written in native platform-dependent code. It loads the classes necessary for the JVM itself, such as those of the `java.lang` package, classes for Java primitives, and so forth. Application classes are loaded using the regular, user-defined class loaders written in Java—so, if needed, the developer can influence the processing of these loaders.

The Class-Loader Hierarchy

The class loaders in the JVM are organized into a tree hierarchy, in which every class loader has a parent. Prior to trying to locate and load a class, a good practice for a class loader is to check whether the class's parent in the hierarchy can load—or already has loaded—the required class. This helps avoid doing double work and loading classes repeatedly. As a rule, the classes of the parent class loader are visible to the children but are not visible otherwise. This structure, which is based on delegation and visibility of the classes, allows for separation of the responsibilities of the class loaders in the hierarchy and makes the class loaders responsible for loading classes from a specific location only.

Let's look at this hierarchy of class loaders in a Java application and explore what classes they typically load. At the root of the hierarchy, Java is the bootstrap class loader. It loads the system classes required to run the JVM itself. You can expect all the classes that were provided with the JDK

The class loaders in the JVM are organized into a tree hierarchy, in which every class loader has a parent. Prior to locating and loading a class, a good practice for a class loader is to check whether the class's parent can load—or already has loaded—the required class.

distribution to be loaded by this class loader.
(A developer can expand the set of classes that the bootstrap class loader will be able to load by using the `-Xbootclasspath` JVM option.)

Note that even though the library might be put on the boot classpath, it won't be automatically loaded and initialized. Classes are loaded into the JVM only on demand, so even though classes might be available for the bootstrap class loader, the application needs to access them to trigger their actual loading. (A curious aspect of this loading process is that you can override JDK classes if your JAR file is prepended to the boot classpath. While this is almost always a poor idea, it does open a door to potentially more-powerful tools.)

A sort of child of the bootstrap class loader is the *extension class loader*, which loads the classes from the extension directories (explained in a moment). These classes may be used to specify machine-specific configuration such as locales, security providers, and such. The locations of the extension directories are specified via the

`java.ext.dirs` system property, which on my machine is set to the following:

```
/Users/shelajev/Library/Java/Extensions:/Library/
Java/JavaVirtualMachines/jdk1.8.0_40.jdk/Contents/
Home/jre/lib/ext:/Library/Java/Extensions:/Network/
Library/Java/Extensions:/System/Library/Java/
Extensions:/usr/lib/java
```

By changing the value of this property, you can change which additional libraries are loaded into the JVM process.

Next comes the system class loader, which loads the application classes and the classes available on the class-

path. Users can specify the classpath using the `-cp` property.

Both the extension class loader and the system class loader are of the `URLClassLoader` type and behave in the same way: delegating to the parent first, and only then finding and resolving the required classes themselves, if need dictates.

The class-loader hierarchy of web applications is a bit more complicated. Because multiple applications can be deployed simultaneously to an application server, they need to be able to distinguish their classes from each other. So, every web application uses its own class loader, which is responsible for loading its libraries. Such isolation ensures that different web applications deployed to a single server can have different versions of the same library without conflicts. So the application server automatically provides every web application with its own class loader, which is responsible for loading the application's libraries. This arrangement works because the web application class loader will try to locate the classes packaged in the application's WAR file first, rather than first delegating the search to the parent class loader.

Finding the Right Class

In general, if multiple classes with the same fully qualified name are available to the JVM, the conflict resolution strategy is simple and straightforward: the first appropriate class wins. The `URLClassLoader`, which most of the class loaders extend from, will traverse the directories in the order they are given on the classpath and load the first class it finds that has requested the class name.

The same goes for JAR files that share the same name. The JAR files will be scanned in the order in which they appear in the classpath, not according to their names. If the first JAR file contains an entry for the required class, the class will be loaded. If not, the classpath scan will continue and reach the

Many security features rely on the class-loader hierarchy for permission checks.

second JAR file. Naturally, if the class isn't found anywhere on the classpath, the `ClassNotFoundException` will be thrown.

Usually, relying on the order of directories in the classpath is a fragile practice, so instead the developer can add the classes to `-Xbootclasspath` to ensure that they will be loaded first. There's nothing in particular wrong with this approach, but maintaining a project that relies on a polluted boot classpath requires work. Intuition about where the classes are loaded from will be broken, and everyone will be confused. A better practice is to resolve the confusion at its root and figure out why there are multiple classes with the same name on the classpath. Maybe upgrading some dependency version, cleaning the caches, or running a clean build will be enough to get rid of the duplicates.

Resolution, Linking, and Verification

After a class is located and its initial in-memory representation created in the JVM process, it is verified, prepared, resolved, and initialized.

- **Verification** makes sure that the class is not corrupted and is structurally correct: its runtime constant pool is valid, the types of variables are correct, and the variables are initialized prior to being accessed. Verification can be turned off by supplying the `-noverify` option. If the JVM process does not run potentially malicious code, strict verification might not be required. Turning off the verification can speed up the startup of the JVM. Another benefit is that some classes, especially those generated on the fly by various tools, can be valid and safe for the JVM but unable to pass the strict verification process. In order to use such tools, the developer should disable this verification, which is often acceptable to do in a development environment.

- **Preparation** of a class involves initializing its static fields to the default values for their respective types. (After preparation, fields of type `int` contain 0, references are null, and so forth.)
- **Resolution** of a class means checking that the symbolic references in the runtime constant pool actually point to valid classes of the required types. The resolution of a symbolic reference triggers loading of the referenced class. According to the JVM specification, this resolution process can be performed lazily, so it is deferred until the class is used.
- **Initialization** expects a prepared and verified class. It runs the class's initializer. During initialization, the static fields are initialized to whatever values are specified in the code. The static initializer method that combines the code from all the static initialization blocks is also run. The initialization process should be run only once for every loaded class, so it is synchronized, especially because the initialization of the class can trigger the initialization of other classes and should be performed with care to avoid deadlocks.

More detail on how the JVM performs the loading, linking, and initializing of classes is explained in [Chapter 5 of the Java Virtual Machine Specification](#).

Other Considerations About Class Loaders

The class-loading model is the central piece of the dynamic operations of the Java platform. Not only does it allow for dynamic location and linking of classes at runtime, but it also provides an interface for various tools to hook into the application.

In addition, many security features rely on the class-loader hierarchy for permission checks. For example, the famous method `sun.misc.Unsafe.getUnsafe()` successfully returns an instance of the `Unsafe` class if it is called from a class that was loaded by the bootstrap class loader. Because only system classes are returned by this loader, every library

that uses the Unsafe API must rely on the Reflection API to read the reference from a private field.

Conclusion

When you're developing a library or a framework, as a rule, you don't have to worry about any issues with class loading. It is a dynamic process that happens at runtime, so you rarely need to influence it. Also, modifying the class-loading scheme rarely benefits a typical Java library.

However, if you create a system of modules or plugins that are intended to be isolated from each other, enhancing the class-loading scheme might be a good idea. Just remember that custom class loaders, being a fundamental force influencing all the classes, can introduce hard-to-spot bugs into literally any part of your application. So take extra care when designing your own class-loading functionality.

In this article, we looked at how the JVM loads classes into the runtime, at the hierarchical model of class loaders Java uses, and the hierarchy model of a typical Java application.

All in all, even if you don't fight class-loading issues or create plugin architectures every day, understanding class loading helps you to understand what is happening in your application. It also provides insight into how several Java tools work. And it really demonstrates the benefits of keeping your classpath clean and up to date. </article>

Oleg Šelajev (@shelajev) is an engineer, author, speaker, lecturer, and developer advocate at ZeroTurnaround. He enjoys spending time tinkering with Clojure, Git, and MacVim and is pursuing a PhD in dynamic software updates and code evolution at the University of Tartu.

LEARN MORE

- [Information on controlling class loaders](#)
- [Class loaders in the JVM Specification](#)



Contexts and Dependency Injection: The New Java EE Toolbox

This series has attempted to demystify Contexts and Dependency Injection (CDI). In the previous articles, which appeared in the last three issues, I discussed what strong typing really means in dependency injection, how to use CDI to integrate third-party frameworks, and how to create loose coupling with interceptors, decorators, and events. This final article covers the integration of CDI with Java EE.

On the web tier, Java EE comes with servlets, WebSockets [See [accompanying article](#). —*Ed.*], and JavaServer Faces (JSF), which are related to the user interface. CDI, whose workings I've explained in the last three articles, can bring the web tier and service tier together to create a homogeneous and integrated application.

Java EE bundles several technologies that enable us to create any kind of architecture, including web application, REST interfaces, batch processing, asynchronous messaging, persistence, and so on. As shown in **Figure 1**, all these applications can be organized in several tiers: presentation, busi-

Java for the service tier. Except for the web client (which uses HTML) and the database (which uses Database Definition Language), most of Java EE uses Java as its primary language, and, therefore, we find Java in most of the application tiers:

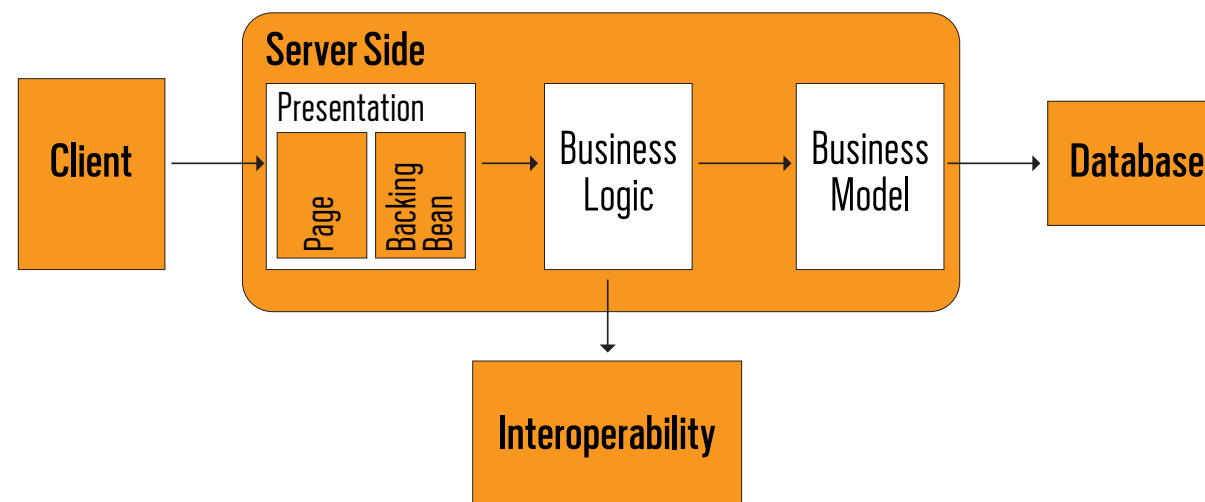


Figure 1. Standard tiers of an application

Java Persistence API entities in the business model or a simple bean on the business logic tier. We even use Java as part of our presentation tier: JSF backing beans are written in Java.

EL for the presentation tier. When I say that Java is the primary language, that's because JSF pages are written using Facelets and Expression Language, or EL. EL provides an important mechanism for enabling the presentation layer to communicate with the application logic. It is used by both JavaServer Faces technology and JavaServer Pages. It uses the # symbol. **Figure 2** shows that EL uses simple expressions to dynamically access data from components—for example, where the purchase order subtotal is displayed on the page or the compute method is invoked when a button is clicked.

CDI to bind service and presentation tiers. To bind both Java and Expression Language, CDI comes to the rescue with a `@Named` annotation. As you can see in **Figure 2**, it basically gives a name to a CDI bean so the bean can be bound in EL. So here, where `PurchaseOrderBean` is annotated with

CDI takes the concept of state management much further, applying it to the entire application, not just to the HTTP layer. Plus, CDI does this in a declarative way: by using a single annotation, the state of the bean is managed by the container.

`@Named("po")`, it means that the bean can be bound in EL with the name `po`.

CDI to manage state. CDI goes further by managing the state of a bean for us using scopes. Let's say that on the top right corner of our web application, we need to display the login of the user. We want this information to remain until the user's session ends. In such a case, we just annotate the bean with `@SessionScoped` and CDI will manage the state by destroying the bean when the session ends. On the other hand, computing and displaying the

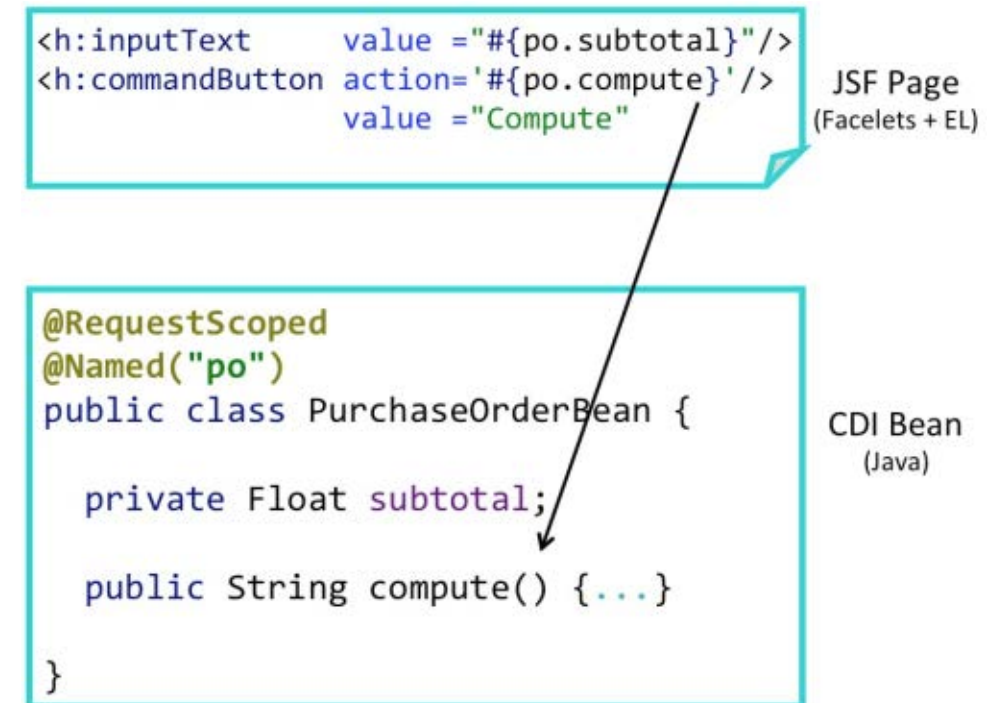


Figure 2. Using Expression Language

total of a purchase order should be done each time the page is refreshed. Because the scope of the `PurchaseOrderBean` must be shorter than the session, we can annotate it with `@RequestScoped`. CDI will maintain the state of the bean only on a per-request basis, which means that this bean is stateless. With just a few annotations, CDI unifies the web tier and service tier, eliminating glue code and letting the developer think about the business problem. CDI defines a uniform model for all our tiers bringing well-defined contexts, which is preserved across multiple requests in a user interaction.

Binding

Binding is the basic service for bringing together the web tier and the service tier. If we want to reference a bean in non-Java code that supports EL, such as a JSF page, we must assign the bean an EL name. The EL name is specified using the `@Named` built-in qualifier. Then we can easily use the bean in any JSF page through an EL expression. EL was orig-

tives can be used. First, we still need to produce, qualify, and name both the VAT and discount rate attributes (see **Listing 5**). Then we add an `@Alternative` annotation. As you can see, CDI is very expressive. Each annotation has its own meaning, and we can read the code very easily. Then it's just a matter of enabling or disabling the alternatives in `beans.xml`.

■ Listing 5.

```
public class NumberProducer {

    @Produces @VAT @Named("vat")
    private Float vatRate = 5.5F;

    @Produces
    @VAT @Named("vat") @Alternative
    private Float vatRateAlt = 19.6F;

    @Produces @Discount @Named("discount")
    private Float discountRate = 2.25f;

    @Produces
    @Discount @Named("discount") @Alternative
    private Float discountRateAlt = 4.75f;

}
```

State Management

We're all used to the concept of an HTTP session and an HTTP request. These are two examples of the broader problem of managing state that is associated with a particular context, while ensuring that all needed cleanup occurs when the context ends—for example, when the HTTP session ends, it needs to be cleaned up. Traditionally, this state management has been implemented manually, by getting and setting servlet session and request attributes. CDI takes the concept of state management much further, applying it to the entire application, not just to the HTTP layer. Plus, CDI does this in a declarative way: by using a single annotation, the state of the bean is managed by the container. No more

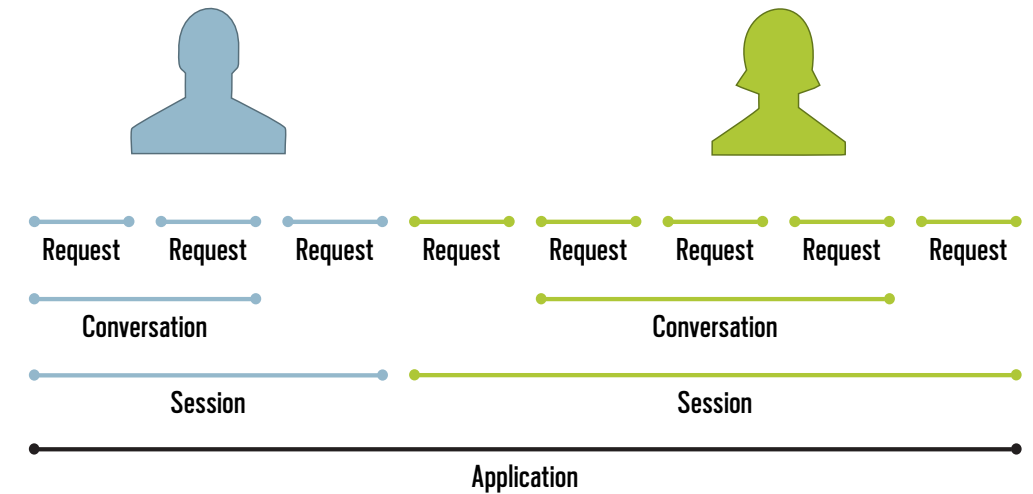



Figure 3. The four built-in CDI scopes

memory leaks when the application fails to clean up session attributes; the CDI container does it automatically. CDI extends the context model defined by the Servlet specification—application, session, request—to another context: a conversation. It then applies the context to the entire business logic, not just to the web tier.

Built-in scope. Before looking at some code, let's first examine the four built-in CDI scopes shown in **Figure 3**. Let's say we have an application that has a lifespan of several months. We boot the server and leave it up and running for a few months before we shut it down. In this case, the application scope lasts for a very long time. One user logs in and remains logged in for a few minutes. The session scope spans from the moment he logs in until the moment he logs out. A second user logs in but her session stays active for a bit longer. Each session is independent and belongs to a single user, and the lifespan can be totally different. In the meantime, both users click at their own pace. Each click creates a request that is handled on the server. The last scope is the conversation and is slightly different because it can span for as long as needed. It's just a matter of beginning a conversation, which can span several requests, and ending it. Each user will have his or her own conversation. Each of these scopes



A language that makes it easy to create projects with libraries from Python and Java

This article gives a detailed tour of Jython 2.7 features, including an easy-to-follow example of working with spreadsheets using Apache POI. After reading this article, you will understand Jython's features enough that you can download the most recent release and get going on your own project.

Running the Jython program without any other arguments starts up the console; we will use the most recently released version as of the writing of this article (2.7.0), as shown in **Listing 1**. (**Note:** In this article, we assume that you are using a UNIX-like system, with \$ being the command-line prompt from a shell such as bash. Jython 2.7 also works well with Windows.)

```
$ jython
Jython 2.7.0 (default:.. . ., Apr 29 2015, 02:25:11)
[Java HotSpot(TM) 64-Bit Server VM (Oracle Corp.)]
```


loads our file. We get the standard prompt of `>>>`. Let's also call the `dir` function to find out what is available; when the function is called with no arguments, it applies to the current module:

```
$ jython27 -i basics.py
>>> dir()
['__builtins__', '__doc__',
 '__file__',
 '__name__', '__package__',
 'inverted']
>>> inverted({1: 'one', 2: 'two'})
{'one': 1, 'two': 2}
```

The `-i` option means we are running the console in the scope of the `basics` module itself; we have a variety of defined names, as well as the `inverted` function we just defined. Using the REPL in this way is very useful for exploratory programming: start the console with the module in progress; try out an idea in the console; edit the module using some extract of this console work; repeat. Such exploratory programming is a hallmark of Python development.

You can also use more-traditional ways of working with your code. So IDEs such as PyDev (built on Eclipse) and PyCharm (built on IntelliJ) support GUI debugging of Python code, including breakpoints, watches, and introspecting variables. This all works because Jython supports Python's standard debugging and tracing mechanisms.

Python 2.7 introduced more functionality when working with the `set` type, which Jython 2.7 supports:

```
>>> {2,4,6,8,10}
set([2, 4, 6, 8, 10])
>>> # Create Empty Set
>>> set()
```

Jython supports
Python's standard
debugging and
tracing mechanisms.

```
set([])
>>> s = {2,4,6,8,10}
>>> 3 not in s
True
```

Given that Jython supports the `set` type, it also supports all of the helpful Python `set` functionality:

```
>>> s.pop()
2
>>> s
set([4, 6, 8, 10])
>>> x.add(3)
>>> x.add(5)
>>> s.symmetric_difference(x)
set([3, 4, 5, 6, 8, 10])
```

Let's now look at Java integration. Jython is not the only way to integrate Python with Java. Other integration options include JPy (embedding CPython via JNI) and Py4J (using a remote socket connection). However, Jython is unique in how it supports working with Java objects as if they are Python objects, and vice versa.

The Python language lacks an ordered set as part of its standard library. But Jython makes it easy to use Java's available implementations of ordered sets, including `java.util.LinkedHashSet`, which maintains insertion ordering, and `java.util.TreeSet`, which maintains natural ordering. **Listing 3** demonstrates how this works.

■ **Listing 3.** (second line wraps)

```
>>> from java.util import TreeSet
>>> clangs = TreeSet(["c", "python",
    "ruby", "perl", "javascript"])
```

One nuance compared with Java is that Python does not use a `new` keyword to construct objects. Instead, the class itself is directly used as a factory, as shown in Listing 4.

■ Listing 4.

```
>>> jvmlangs = TreeSet(["java", "python",
"groovy", "scala", "ruby", "javascript"])

>>> clangs | jvmlangs # set union
[c, groovy, java, javascript, perl,
python, ruby, scala]

>>> clangs & jvmlangs # set intersection
[javascript, python, ruby]
```

Let's now try working with other Java packages. The developers of Jython are fans of the collections support in the Google Guava library, using it extensively—especially MapMaker, the library's concurrent map.

First, we need to download the Google Guava JAR file and put it in our `CLASSPATH`. (In this example, we're using [Guava release 18.0.](#)) Restart the Jython console with the `CLASSPATH` change.

Now is a good time to try out the tab completion support newly available in Jython 2.7. Sometimes, working with the Java ecosystem is a bit unwieldy: the package names are deep and often spelled out. With tab-completion support, at any point as you're typing in the console, you can press the TAB key to get a list of possible completions (if there are multiple possibilities) or the completion itself (if there is just one possibility).

So first import:

```
>>> import com.google
```

Then you can type, followed by pressing the TAB key:

```
>>> d = com.google.c
```

And you will get the following:

```
>>> d = com.google.common
```

You can eventually complete it like this:

```
>>> d = com.google.common.collect.HashBiMap.create(
... dict(one=1, two=2, three=3))
```

The advantage of this bidirectional mapping is that it is maintained over any updates, as we see here:

```
>>> d.inverse()
{3: three, 2: two, 1: one}
>>> d.inverse()[4] = "four"
>>> d
{three: 3, four: 4, two: 2, one: 1}
```

Project: Working with Spreadsheets

We are now going to explore an in-depth example to demonstrate Jython's deep support for Java integration.

Our premise is that of automating an existing business process that relies on consolidating spreadsheets to report on the state of the business. Although this spreadsheet-based process has proven its flexibility, it is also manual and error-prone. Currently this process relies on e-mail, shared drives, macros, and some proprietary tools. Sound familiar?

As software developers, we know that there are many ways this business process could be automated. We could rewrite it to remove the use of spreadsheets altogether. But we want to preserve the advantages of spreadsheets—including their wide adoption, flexibility, and ease of use. So let's try a different approach: we will continue to work with spreadsheets, but provide better tooling for managing them. We will use and integrate the following:

- [Apache POI](#), to programmatically work with spreadsheet workbooks (Java library).
- [GitHub](#), to version workbooks; we especially want to take advantage of its extensive REST API to store and retrieve workbooks (REST service). GitHub is, therefore, representative of general REST services we might use to manage



workbooks, including ones we might build ourselves.

- Requests, to simplify using HTTP and especially RESTful services to use GitHub's support of a REST API to resolve artifacts (Python library).
- Nosetests, to support unit testing (Python library).
- Custom Python code to glue together all of the above, including auditing and formula computation.

First, download Apache POI. As of the writing of this article, the latest version is 3.12. You will need to add JAR files from the POI distribution that support poi, poi-ooxml, poi-ooxml-schemas, and xmlbeans to your **CLASSPATH**.

Next, you need to install the required Python packages. Installation support is a highlight of the Jython 2.7 release. In the past, for Python developers, one of the aggravating things about earlier releases of Jython was that they never contained full support for the Python ecosystem. Now with version 2.7, you can readily take advantage of the Python ecosystem, specifically the large number of Python packages that are available on PyPI (Python Package Index). Doing so is quite easy in Jython 2.7, because the popular pip tool is bundled with the release. This support for tooling makes it easy to incorporate the libraries and APIs of the Python ecosystem into an application.

The following command will install Nostests and Requests modules by using the pip module. The `-m MODULE` means run the desired Python module as if it were a command-line program, passing the remaining arguments to it:

```
$ jython27 -m pip install nose requests
```

What makes the Python language so nice is that we can incrementally explore in a console both the problem space and possible solutions.

With both Java and Python dependencies now taken care of, where should we start? What makes the Python language so nice is that we can incrementally explore in a console both the problem space and possible solutions.

Assume that we have a simple spreadsheet named `hours.xlsx` located at the top level of the GitHub repo, <https://github.com/jimbaker/poi>. We can retrieve this spreadsheet as <https://github.com/jimbaker/poi/raw/master/hours.xlsx>. Let's try this from the console, where `url` is set to the desired spreadsheet:

```
>>> import requests
>>> response = requests.get
... (url, stream=True)
```

Let's write the response to a binary file (so the file mode is `"wb"`); we will do so iteratively in 512-byte chunks so that memory consumption is minimized. The `writelines` method takes an iterator:

```
>>> f = open("hours.xlsx", "wb")
>>> f.writelines(response.iter_content(512))
>>> f.close()
```

Now let's read the saved spreadsheet with POI; note that Jython implicitly bridges Python file objects as `FileInputStream` or `FileOutputStream` to use Java methods or constructors, as needed:

```
>>> from org.apache.poi.xssf.usermodel
... import XSSFWorkbook
>>> workbook = XSSFWorkbook(
... open("hours.xlsx", "rb"))
```

What can we do with this open workbook? Let's explore:

```
>>> dir(workbook)
```


Eventually after some exploration in the console and the POI API documentation, we might arrive at code like **Listing 5** to process a workbook.

■ Listing 5.

```
# traverses cells in a workbook,
# calling a callback function on each cell

from contextlib import closing

def process_workbook(path, callback=None):
    if callback is None:
        def callback(cell):
            print cell,

    with open(path, "rb") as file:
        with closing(XSSFWorkbook(file)) as workbook:
            for sheet in workbook:
                for row in sheet:
                    for cell in row:
                        callback(cell)
```

The `process_workbook` function takes two parameters, `path` and `callback`. Note that `callback` is an optional parameter, because we give it a default value of `None`. However, there is no statically specified type as we would see in Java or would possibly be inferred in Scala. When we talk about static analysis of a program (or lexical analysis), it means that by only examining the program text, not running it, we—or a compiler or an IDE or some other tool—can determine certain properties of the program. What is the scope of the variables? What are the types of variables? Are types used in a consistent fashion—in other words, does the code type check? Is a chunk of code unreachable (or dead) and, therefore, can be eliminated? Can we use constant folding or inlining? And so forth. Of the items mentioned here, Jython supports only determining variable scope statically. (CPython does some amount of constant folding and dead code elimi-

nation, and Python 3.5 will have standard static type annotations as part of its support for gradual typing, a type system that mixes both dynamic and static type approaches.)

We also define a function if `callback` is not defined, which we call, perhaps confusingly at first, `callback`. This function is inside the scope of the `process_workbook` function. But not only is it lexically scoped—and is, in fact, a closure—the `callback` function is conditionally defined. So that's really quite different from what we would do in Java. Once again, Python is showing that it is certainly a dynamic language; any static analysis of `process_workbook` could conclude only that `callback` might be this function. Or it might not be. However, do note that Jython has already compiled the source code to Java bytecode. So what we are seeing here is whether the name `callback` will be set to the corresponding compiled function body. Consequently the overhead of this conditional definition is actually no more than the overhead for some variable assignments. This shows how Jython enables you to move back and forth between the Java and Python ways of doing things.

We then iterate over the spreadsheets in each workbook, the rows in each spreadsheet, and the cells in each row. The workbook, spreadsheet, and row objects all implement `java.lang.Iterable`, which Jython can iterate over. Perhaps not surprisingly, Jython’s integration ensures that Java code can also use for-each loops to iterate over Python iterables (and iterators).

When `callback` is called against `cell` with `callback (cell)`, the Jython runtime does dynamic type checking: is the `callback` object a callable? Python has a simple rule: all callable objects implement a specific method, `__call__`.

Jython enables you to move back and forth
between the Java
and Python ways
of doing things.


```
def get_cells(sheet, ref):
    for row_idx, col_idx in referred(ref):
        row = sheet.getRow(row_idx)
        if row is not None:
            cell = row.getCell(col_idx)
            if cell is not None:
                yield cell
```

The `get_cells` function is a *generator* function: calling this function returns an iterator that successively yields values on each iteration (as marked by the `yield` keyword). So this is a very convenient way of constructing the equivalent of `java.lang.Iterator`, but without having to explicitly capture state between each invocation of the `next` method. Using generators is commonly seen in idiomatic Python code, especially because doing so simplifies incrementally working with data—particularly large data sets. **Listing 8** (available in the [download area](#)) demonstrates using generators.

With `get_cells`, we can quickly compute answers to a number of queries. Let's try it out. What is the sum of the range of A1:G8? In other words, what is equivalent to `=SUM(A1:G8)` in the spreadsheet?

Define the following helper function `get_nums` and use the built-in function `sum`:

```
NUMERIC_CELLS = {
    Cell.CELL_TYPE_FORMULA,
    Cell.CELL_TYPE_NUMERIC }
```

```
def get_nums(cells):
    for cell in cells:
        if cell.cellType in NUMERIC_CELLS:
            yield cell.numericCellValue
```

Using the built-in function `sum`, the answer becomes simply `sum(get_nums(get_cells(spreadsheet, "A1:G8")))`.

Are any cells in the range A1:G8 using hardcoded formulas?
Define a variant of the hardcoded audit function we had

earlier and use the built-in function `any`:

```
def hardcoded_cells(cells):
    for cell in cells:
        try:
            float(cell.cellFormula)
            yield True
        except:
            yield False
```

The answer is `any(hardcoded_cells(get_cells(spreadsheet, "A1:G8")))`.

What we have done here is define the beginning of a high-level Python API for working with workbooks that resembles, in part, the functions that we might use in the spreadsheet itself. However, we also retain all the low-level Java POI library functionality as well.

This leads us to the last topic: are we able to ensure that our spreadsheets pass a set of compliance tests? This is a bit more involved, but let's say we set up a continuous integration service such as Jenkins to run tests on spreadsheets, perhaps as part of a GitHub pull request. How do we define and run our tests? The Python ecosystem has several good choices for testing frameworks, including in the standard library itself and `unittest`, which is an implementation of the xUnit style of testing. But the Nose testing framework, which builds on `unittest`, is justifiably popular because it is easy to use.

For example, we might want to ensure the correctness of our cross-tabulations: the sum of subtotals along a row should equal the sum of subtotals along a column, taking into account numerical precision issues, as shown in **Listing 9**.

■ **Listing 9.**

```
from nose.tools import assert_almost_equals
def assert_crosstab(spreadsheet, range1, range2):
    assert_almost_equals(
        sum(get_nums(spreadsheet, range1)),
        sum(get_nums(spreadsheet, range2)))
```


Once this is defined, you can write a simple test script like Listing 10.

■ **Listing 10.**

```
finance_wb = XSSFWorkbook("financials.xlsx")
main_sheet = finance_wb.sheetAt(0)
```

```
def test_financials():
    assert crosstab(main sheet, "A5:G5", "H1:H5")
```

Then you can run Nose to discover and run your tests. Nose follows a convention-over-configuration approach, which means it is easy to get going. The result is:

```
$jython -m nose
.
-----
Ran 1 test in 0.033s
```

OK

Each dot on the second line corresponds to a test in all of the collected test files. Simply add more tests.

Don't Look Back!

Predictably, as time marches on, technologies and language features evolve. As such, there have been some important deprecations in Jython 2.7. Perhaps the most notable is that Jython 2.7 requires a minimum of Java 7. Another important note is that the installer no longer supports the use of an alternative JRE when generating Jython launchers. Use `JAVA_HOME` instead.

Jython 3.5

The Python language undergoes continued active development, along with its reference implementation. By the time you read this article, CPython 3.5 will be released. At some

future point, a release of Jython 3.5 is planned, paralleling the CPython 3.5 release. It is worth pointing out that Jython 2.7 basically has the same internal runtime support and stdlib as Python 3.2. But substantial work will be required to get to Jython 3.5. One eagerly anticipated feature in Python 3.5 is optional static typing, which will enable even better Java integration in Jython.

But not so fast. Jython 2.7.x will be around for quite some time. The Jython team plans to work on 2.7.x as long as Python 2.7 is in wide use. The migration to and adoption of Python 3 has been fairly slow, partly due to the significant changes between versions 2.7 and 3.0. Because Python 2.7 is still in wide use, the Jython team plans to make time-based releases of the Jython 2.7.x line. The future releases of Jython 2.7.x will focus on performance, integration, and more. The release of Java 9 will likely improve performance with more optimization for dynamic languages.

Although there is no rush to get to Jython 3.5, it is on the docket. In fact, there is a branch of development already devoted to Jython 3.5, although it is in the early stages. Currently the target for a Jython 3.5 release is in the next two years.

Conclusion

Jython 2.7 provides a wealth of tools, enabling developers to combine two of the most popular ecosystems, Python and Java, in the same codebase. In this article, we took a look at some of the top new features of Jython 2.7, but there are plenty of other great features to explore. Download it from jython.org and watch for updates, which are planned for every six months. [</article>](#)

[This article is part of an ongoing series exploring JVM languages. In the last issue, we covered Kotlin. In the next issue, we examine Gosu, a JVM language used in industry for both front ends and back-end systems. —Ed.]

The lightweight virtualization container is fast becoming the preferred way to package and deploy Java web apps.



An application typically requires a specific version of the operating system, JDK, application server, database server, and a few other infrastructure components. In order to provide an optimal experience, it might need binding to specific

- Images: the build component of Docker, which consists of a read-only template of the application operating system. For example, an image could be the Fedora operating system with WildFly and your Java EE application installed. You can

Docker Toolbox is the easiest way to get started with Docker.

- **Images:** can be used to easily create new images or update existing ones.
- **Containers:** a runtime representation created from images. Containers are the run component of Docker. They can be run, started, stopped, moved, and deleted. Each container is an isolated and secure application platform.
- **Registry:** the distribution component of Docker, where images are uploaded and downloaded. A registry can be public, such as [Docker's own registry](#). A private registry inside your firewall can be set up as well.

How Does Docker Work?

Docker uses a client/server architecture. The Docker daemon is the server and runs on a host machine, where it does the heavy lifting of running Docker containers. The Docker client, a binary that can be installed on your machine, communicates with the Docker daemon and sends administrative commands such as a command to retrieve an image or run a container. The Docker registry is where all the images are stored. **Figure 1** shows this typical layout.

The Docker host may be colocated with the Docker client in early development stages. But it's generally moved to a separate machine for scalability reasons.

A typical workflow entails the following:

- The Docker client asks the Docker daemon to run the container for a given image.

- The Docker daemon checks whether the image already exists on the host. If it does, then the daemon runs the container. If not, then it downloads the image from the Docker registry and runs the container. Multiple containers that use the same image can be run easily.

The Docker client has commands for building and updating images; downloading and uploading images to a registry; running, querying, watching, and killing the running containers; and much more. It communicates with the Docker daemon using a socket or REST API. The Docker daemon talks to the Docker registry, if required, to perform the needed operation.

As a developer, you build the image for your application, run containers using that image, and then upload that image to the Docker registry for others to try it.

Getting Started with Docker

Linux machines natively support Docker, and it's easily installed using the default package manager. If you're using a Windows or Mac system, [Docker Toolbox](#) provides the different tools required to get started with Docker. It contains

- The Docker client (**docker** binary). This is the same client component previously discussed. It's used to manipulate images and containers by communicating with the Docker daemon.
- Docker Machine (**docker-machine** binary). Docker Machine lets you create Docker hosts on VMs that reside on your computer, with cloud providers, and inside your own data center. The Docker daemon is then installed inside this VM, after which a client can be configured to talk to this host.

Docker Machine creates a VM using *drivers*. A driver is an overloaded term that here means a virtual environment. For example, the Oracle VM VirtualBox driver is used on a local Mac or Windows system. AWS, Microsoft Azure, and other drivers are available to create a host in the cloud.

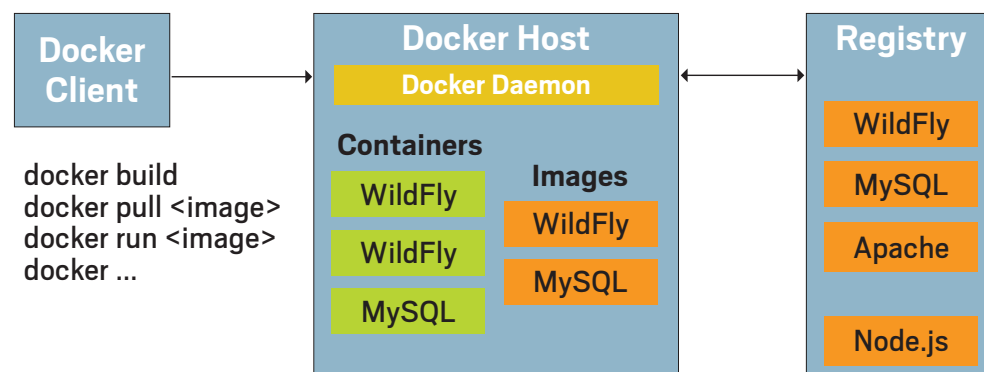


Figure 1. A typical Docker setup

`docker ps` shows the list of running containers. Presently, the list of containers is empty because no containers have been started. If you also want to include previously terminated containers, then the `-a` option needs to be specified as an option. The output of this command is best seen with a width of 128 columns.

Note that `docker --help` shows the complete list of commands. Similarly, `docker ps --help` shows all the options available for a command. Feel free to check out different commands to see how you can modify them to meet your needs.

Now, let's run the prebuilt "Hello World" Docker image, which is found on Docker Hub. This is done using the following command:

```
docker run hello-world
```

The command shows the following output (some more verbiage is shown before and after this text in the output):

Hello from Docker.

The output verifies that

- The Docker client and daemon are installed correctly.
- The hello-world image is not available on the Docker daemon, but Docker was able to download the image from Docker Hub.
- The container runs from the downloaded image and streams the output to the Docker client.

As you can see, you can run your first container without much fuss.

Build Your First Docker Image Using Java

Docker images are read-only templates that launch Docker containers, and each image consists of a series of layers. Docker makes use of a *union file system* to combine these

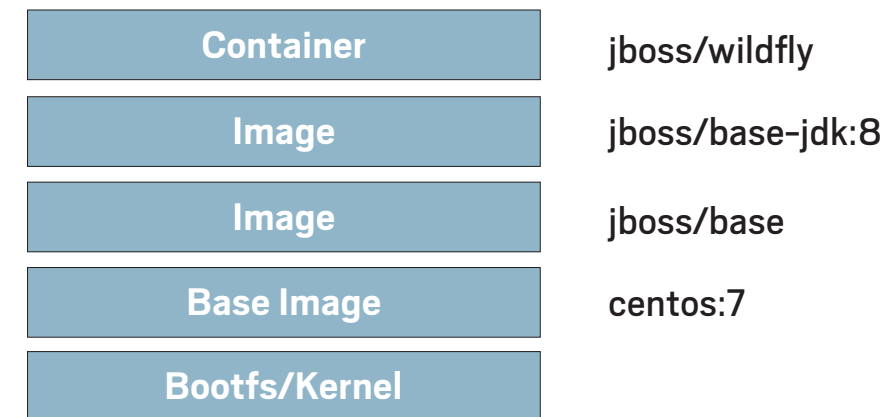


Figure 2. Building a Docker image for Java EE

layers into a single image. Union file systems enable files and directories of separate file systems to be transparently overlaid, forming a single, coherent file system.

This layering makes Docker extremely lightweight. Any change to a Docker image—for example, updating an application to a new version or changing the JDBC driver—requires only that the affected layer be rebuilt. Thus, rather than replacing the whole image or entirely rebuilding, as you do with a VM, only one layer is added or updated. This also makes distribution faster and simpler.

Every image starts from a base operating system image. For example, `fedora` is a base Fedora image. Multiple layers are then added. For example, `jboss/wildfly` is built using multiple images, as shown in Figure 2.

A Docker image is built by reading the instructions from a text file, which is usually called **Dockerfile**. This file contains all the commands needed to build the image. For example, it specifies the base operating system, the JDK, the application server, and any other dependencies. The JDK and application server are downloaded and installed in the image using typical shell-like commands, such as **GET**, **COPY**, and **RUN**. The **COPY** instruction can be used to copy files from the local file system into the container as well. Optionally, you might consider using a base image that already contains the JDK or WildFly and build upon that. Docker Hub has a base


```
docker inspect --format '{{.Config.ExposedPorts }}'
<CONTAINER_ID>
```

Furthermore, the `-P` option can be used to map the container port to a high port (that typically is in the range of 23768 to 61000) on the local host. Then, container port-to-host port mapping can be seen:

```
docker port <CONTAINER_ID>
8080/tcp -> 0.0.0.0:32768
```

In this case, the application will be accessible at `dockerhost:32768/movieplex7`.

Conclusion

This article explained the key concepts of Docker and how to package your Java applications using it. Docker enables the Package Once Deploy Anywhere (PODA) paradigm, and it is changing how applications are built, deployed, and scaled. Docker reduces the impedance mismatch between development, testing, and production environments.

Ready or not, Docker is here and likely to be with us for a long time as a lightweight container technology. In the next article in this series, I'll discuss multicontainer apps and running containers in clusters. `</article>`

LEARN MORE

- [Getting started with Docker](#)
- [Overview of containers](#)
- Kubernetes: Docker orchestration tool

PUNE JAVA USER GROUP



Pune, with more than three million residents, is the ninth most populous city in India. It is also the second-highest software exporting city in India. Pune has been a center of learning and academia over many centuries and boasts several prestigious educational institutions.

The Pune Java User Group is one of the oldest in India. It was founded in the late 1990s and has been active ever since. The Java user group (JUG) leadership has changed several times over the years, but the focus has always been on learning new Java technologies and encouraging discussions among Java enthusiasts.

The Java language and Java EE are the usual topics of discussion. However, the group also welcomes discussions of other languages that run on the Java platform. For example, a recent talk at the JUG featured Java Champion Andres Almiray discussing the Groovy language and the build tool Gradle.

The JUG meetings are held at educational institutions and software companies in Pune. The group has more than 1,400 members and is active via multiple social media channels, such as Twitter ([@JavaPune](#)) and [Google Groups](#). The launch events for new versions of Java have had hundreds of enthusiasts in attendance.

Pune has a buzzing startup ecosystem, and the JUG has served as a platform for interactions among tech startups. The JUG has also assisted in building tech communities in the city around many other niche as well as mainstream technologies.

The easy-to-use API for long-lived connections



For many kinds of web applications, having the user always in the driver's seat is not desirable. From financial applications with live market data, to auction applications where people around the world bid on items, to the lowly chat and presence applications, web

applications have long sought means by which the server side can push data out to the client. A mix of ad hoc mechanisms arose out of this need that were either based around keeping long-lived HTTP connections or some form of client polling; none proved a complete solution to the problem. The need for a new approach led to the development of the WebSocket protocol.

The WebSocket protocol is a TCP-based protocol that provides a full duplex communication channel over a single connection. In simple terms, this means that it uses the same underlying network protocol as does HTTP and that over a single WebSocket connection both parties can send messages to the other at the same time. The WebSocket protocol defines a simple connection lifecycle and a data-framing mechanism that supports binary and text-based messages. Unlike HTTP, the connections are long lived. This means that because the connection need not continually be re-established for each message transmission, as the anti-symmetric HTTP protocol does, every data message in the WebSocket protocol does not need to carry all the meta-information about the connection, as HTTP does. In other words, once the connection is established, the message transmission is much lighter weight than in the HTTP protocol.

However, this is not the primary reason WebSocket is better suited to the task of servers pushing information than are polling frameworks layered on top of the HTTP proto-


```
Sec-WebSocket-Accept: HSmrc0sM1YUkAGmm50PpG2HaGwk=  
Sec-WebSocket-Protocol: chat  
Sec-WebSocket-Extensions: compress, mux
```

This response confirms that the server will accept the incoming TCP connection request from the client and may impose restrictions on how the connection may be used. Once the client has processed the response and is happy to accept any such restrictions, the TCP connection is created, as shown in **Figure 1**, and each end of the connection may proceed to send messages to the other.

Once the connection is established, several things can occur:

- Either end of the connection may send a message to the other. This may occur at any time that the connection is open. Messages in the WebSocket protocol have two flavors: text and binary.
- An error may be generated on the connection. In this case, assuming the error did not cause the connection to break, both ends of the connection are informed. Such nonterminal errors may occur, for example, if one party in the conversation sends a badly formed message.
- The connection is voluntarily closed. This means that either end of the connection decides that the conversation is over and so closes the connection. Before the connection is closed, the other end of the connection is informed of this.

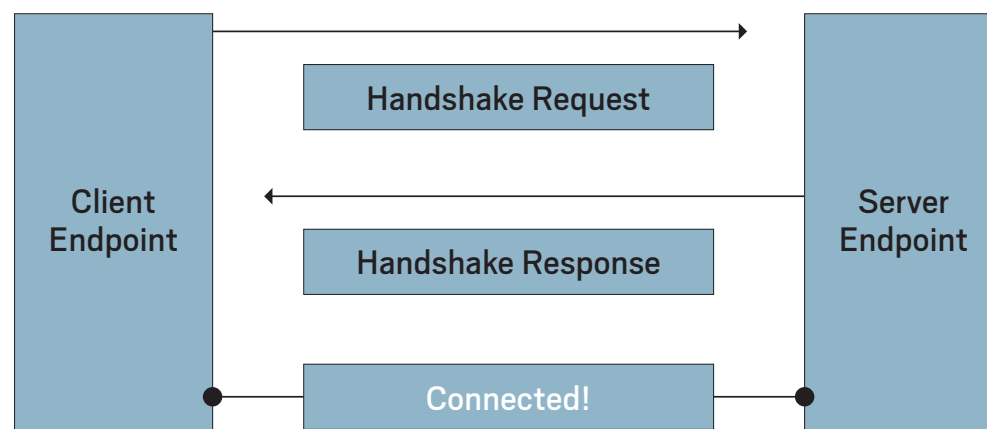


Figure 1. Establishing a WebSocket connection

Overview of the Java WebSocket API

The Java WebSocket API provides a set of Java API classes and Java annotations that make it relatively straightforward to create WebSocket endpoints that reside in the Java EE web container. The general idea is to take a Java class in which you want to implement the logic of the server endpoint and annotate it at the class level with the special Java WebSocket API annotation `@ServerEndpoint`. Next, you annotate its method with one of the lifecycle annotations, such as `@OnMessage`, which imbues the method in question with the special power of being called every time a WebSocket client sends a message to the endpoint. You then package it in the `WEB-INF/classes` directory of the WAR file. Listing 1 is an example of this.

■ Listing 1. The EchoServer sample

```
import javax.websocket.OnMessage;
import javax.websocket.server.ServerEndpoint;

@ServerEndpoint("/echo")
public class EchoServer {

    @OnMessage
    public String echo (String incomingMessage){
        return "I got this (" +
                incomingMessage + ")" +
                " so I am sending it back !";
    }
}
```

This WebSocket endpoint is mapped to `/echo` in the URI space of the web application. Each time a WebSocket client sends it a message, it responds back immediately with a message derived from the one it received.

modes. It allows you to translate WebSocket messages to and from arbitrary Java classes using decoder and encoder classes.

The Java WebSocket API also provides the means to create WebSocket client endpoints. The only time that the WebSocket protocol is asymmetric concerns who initiates the connection. The client support in the Java WebSocket API allows a client to connect to the server, and so is suitable for Java clients to connect to WebSocket endpoints running in the Java EE web container or, in fact, any WebSocket server endpoint.

Before we look at a real example of a Java WebSocket, let's take a tour of the annotations and main classes in the Java WebSocket API. Don't worry about spending too long before we get to working code: the Java WebSocket API is one of the smaller APIs of the Java EE platform.

WebSocket Annotations

The Java WebSocket annotations have two main purposes. First, they allow you to declare that you want an ordinary

Java class to become a WebSocket endpoint, and second, they allow you to annotate methods on that class so that they intercept the lifecycle events of the WebSocket endpoint. First, we will take a look at the class-level annotations.

@ServerEndpoint. This is the workhorse annotation of the API, and if you create many WebSocket endpoints, you will be seeing a lot of it. The only mandatory attribute of this class-level annotation is the value attribute (see Table 1), which specifies the URI path to which you want this endpoint to be registered in the URI space of the web application.

@ClientEndpoint. The [@ClientEndpoint](#) annotation is used at the class level on a Java class that you wish to turn into a client endpoint that initiates connections to server endpoints. This is often used in rich client applications that connect to the Java EE web container. It has no mandatory attributes.

@ServerEndpoint and @ClientEndpoint optional attributes. As shown in Table 2, these class-level annotations have several other attributes in common that define other configuration

options that apply to the WebSocket endpoint they decorate.

Now let us turn to the lifecycle annotations.

@OnOpen. This method-level annotation declares that the Java EE web container must call the method it annotates on a WebSocket endpoint whenever a new party connects to it. The method may have either no arguments or an optional [Session](#) parameter, where the class [javax.websocket.Session](#) is an API object that represents the WebSocket connection that has just opened; and/or an optional Endpoint config parameter, where [javax.websocket.EndpointConfig](#) is an API object representing the con-

ATTRIBUTE	FUNCTION	MANDATORY
VALUE	DEFINES URI PATH UNDER WHICH THE ENDPOINT IS REGISTERED	YES

Table 1. The attribute of [@ServerEndpoint](#)

@SERVERENDPOINT AND @CLIENTENDPOINT ATTRIBUTES	FUNCTION	MANDATORY
CONFIGURATOR	THE CLASS NAME OF A SPECIAL CLASS THE DEVELOPER MAY PROVIDE TO DYNAMICALLY CONFIGURE THE ENDPOINT	NO
DECODERS	LIST OF CLASSES USED TO CONVERT INCOMING WEBSOCKET MESSAGES INTO JAVA CLASSES THAT REPRESENT THEM	NO
ENCODERS	LIST OF CLASSES USED TO CONVERT JAVA CLASSES INTO OUTGOING WEBSOCKET MESSAGES	NO
SUBPROTOCOLS	LIST OF STRING NAMES DENOTING ANY SPECIAL SUBPROTOCOLS, SUCH AS "CHAT," THAT THE ENDPOINT SUPPORTS	NO

Table 2. Attributes of class-level annotations

WebSocketContainer. As the `ServletContext` is to Java servlets, so is the `WebSocketContainer` to Java WebSockets. It represents the Java EE web container to the WebSocket endpoints it contains. It holds a number of configuration properties of the WebSocket functionality, such as message buffer sizes and asynchronous send timeouts.

Let’s Build Something: A WebSocket Clock

We have completed our tour of the Java WebSocket API, and having done so, we know more than enough to look at our first WebSocket application. The Clock application is a simple web application. When you run the application, you see the `index.html` web page, as shown in Figure 2.

When you click the Start button, the clock starts with the current time, as shown in Figure 3. The time updates every second.




Figure 2. WebSocket Clock off




Figure 3. WebSocket Clock on

When you click the Stop button, the clock stops until you restart it, as shown in Figure 4.

The application is made up of a single web page, `index.html`, and a single Java WebSocket endpoint, called `ClockServer`. When Start is pressed, `index.html` uses JavaScript code to establish a WebSocket connection with the `ClockServer` endpoint. It sends time update messages every second back to the browser client. The JavaScript code handles the incoming message and renders it on the page. Clicking Stop causes the JavaScript code in the `index.html` page to send a stop message to the `ClockServer`, which consequently stops sending the time updates. This architecture is shown in Figure 5.

Let’s look at the code, first for the client. [The complete listing is available for download at this issue’s [download area](#). —Ed.] Here is the WebSocket client code from Listing 2.




Figure 4. WebSocket Clock stopped




Figure 5. Clock architecture


```

        return "unknown message: " +
                incomingMessage;
    }
}

@OnError
public void clockError(Throwable t) {
    this.stopClock() ;
}

@OnClose
public void stopClock() {
    this.running = false;
    this.updateThread = null;
}
}

```

Notice that the `ClockServer` uses the `@ServerEndpoint` annotation to declare itself as a WebSocket endpoint, mapped to the URI `/clock`, relative to the context root of the web application in which it is contained. Notice that the `startClock()` method, called when a new client connects thanks to its `@OnOpen` annotation, does most of the work. It creates a thread that uses the `Session` object to obtain a reference to the `RemoteEndpoint` instance representing the client and sends it the current time, formatted into a string. If the endpoint receives a message, it is passed into the `handleMessage()` method, which you can identify because this method is annotated with `@OnMessage`. The `String` parameter of this method informs you that the endpoint is electing to receive its text messages in the simplest form of a Java string. This method returns a string, which is turned into a WebSocket message by the Java EE container and sent back to the client immediately.

How Many WebSocket Instances?

One question that arises even in this simple example is: how many instances will occur for a WebSocket endpoint class

such as the `ClockServer`? The answer is that there will be one instance of the `WebSocket` endpoint class for each client that connects to it. Each client gets a unique endpoint instance. Further, the Java EE web container guarantees that no two `WebSockets` are sent to the same endpoint instance at once. So, in contrast to the Java servlet model, you can program your `WebSocket` endpoints knowing that there will only ever be one thread calling at a time.

Conclusion

The base WebSocket protocol gives us two native formats to work with: text and binary. This works well for very simple applications that exchange only simple information between client and server. For example, in our Clock application, the only data that is exchanged during the WebSocket messaging interaction is the formatted time string broadcast from the server endpoint and the `stop` string sent by the client to end the updates. But as soon as an application has anything more complicated to send or receive over a WebSocket connection, it will find itself seeking a structure into which to put the information. As Java developers, we are used to dealing with application data in the form of objects: either from classes from the standard Java APIs, or from Java classes that we create ourselves. This means that if you stick with the lowest-level messaging facilities of the Java WebSocket API and want to program using objects that are not strings or byte arrays for your messages, you need to write code that converts your objects into either strings or byte arrays and vice versa. I'll discuss this topic in the second installment of this article. [</article>](#)

This article was adapted from the book [Java EE 7: The Big Picture](#) with kind permission from the publisher, Oracle Press. The book was reviewed on page 10 of the September/October issue.

LEARN MORE

- Oracle's Java WebSockets tutorial

Quiz Yourself

Easy questions are hard for the wise man. And hard questions are easy. But what is the place for wisdom in quizzes? Let's see.

The questions in this quiz section are taken from certification test 1Z0-808: Oracle Certified Associate, Java SE 8 Programmer, [Oracle Certified Java Programmer](#). More than last issue's questions, these little questions have unexpected traps that can snag even the attentive coder. Many times you can avoid snags by coding predictably down the middle of the language. But as we see here, even then not everything works out exactly as we expect—more so, because here we dip into a few streams, which have traps all their own. (Answers appear in the “Answers” section immediately after the questions.)

Question 1. Given these code fragments:

```
class ProductNotFoundException extends Exception{
```

```
class SalesPerson {
    String name;
    List<String> products = new ArrayList<>();
    public List<String> getProducts() throws
        ProductNotFoundException {
        products.add("SoundCard");
        return products;
    }
}
```

and

```
class SalesApp {
    public static void main(String[] args) {
```

```

        SalesPerson sp = new SalesPerson();
        List<String> products = sp.getProducts();
        System.out.println(products.get(0));
    }
}

```

What is the result?

- a. SoundCard
- b. A `ProductNotFoundException` is thrown at runtime.
- c. 0
- d. A compilation error occurs.

Question 2. Given this code fragment:

```
String wishMsg = "Happy day!";
wishMsg.concat(" Tom");
String msg = (wishMsg.length() > 10) ?
    "Too long" : "Sent";
System.out.println(msg+": "+wishMsg);
```

What is the result?

- a. Sent: Happy day!
- b. Too long: Happy day!
- c. Too long: Happy day! Tom
- d. A compilation error occurs.

Question 3. Given the content of the AClass.java, BClass.java, and IFace.java files:

```
public abstract class AClass {
    public void aMethod() {
```




Article Proposals

source or those bundled with the JDK). Finally, algorithms, unusual but useful programming techniques, and most other topics that hard-core Java programmers would enjoy are of great interest to us, too. Please contact us with your ideas at javamag_us@oracle.com and we'll give you our thoughts on the topic and send you our nifty writer guidelines, which will give you more information on preparing an article.

Customer Service

+1.847.763.9635), who will do whatever they can to help.

Where?

While it will have no influence on our decision whether to publish your article or letter, cookies and edible treats will be gratefully accepted by our staff at *Java Magazine*, Oracle Corporation, 500 Oracle Parkway, MS OPL 3A, Redwood Shores, CA 94065, USA.

