

OSGi In Practice

Neil Bartlett

January 11, 2009

Contents

Preface	xiii
I Nuts and Bolts	1
1 Introduction	3
1.1 What is a Module?	4
1.2 The Problem(s) with JARs	5
1.2.1 Class Loading and the Global Classpath	6
1.2.2 Conflicting Classes	8
1.2.3 Lack of Explicit Dependencies	9
1.2.4 Lack of Version Information	10
1.2.5 Lack of Information Hiding Across JARs	12
1.2.6 Recap: JARs Are Not Modules	12
1.3 J2EE Class Loading	13
1.4 OSGi: A Simple Idea	15
1.4.1 From Trees to Graphs	16
1.4.2 Information Hiding in OSGi Bundles	18
1.4.3 Versioning and Side-by-Side Versions	19
1.5 Dynamic Modules	19
1.6 The OSGi Alliance and Standards	20
1.7 OSGi Implementations	21
1.8 Alternatives to OSGi	21
1.8.1 Build Tools: Maven and Ivy	22
1.8.2 Eclipse Plug-in System	22
1.8.3 JSR 277	23
2 First Steps in OSGi	25
2.1 Bundle Construction	25
2.2 OSGi Development Tools	26
2.2.1 Eclipse Plug-in Development Environment	26
2.2.2 Bnd	27
2.3 Installing a Framework	28
2.4 Setting up Eclipse	29
2.5 Running Felix	31
2.6 Installing bnd	33
2.7 Hello, World!	34

2.8	Bundle Lifecycle	36
2.9	Incremental Development	39
2.10	Interacting with the Framework	40
2.11	Starting and Stopping Threads	43
2.12	Manipulating Bundles	43
2.13	Exercises	44
3	Bundle Dependencies	47
3.1	Introducing the Example Application	48
3.2	Defining an API	48
3.3	Exporting the API	51
3.4	Importing the API	53
3.5	Interlude: How Bnd Works	57
3.6	Requiring a Bundle	59
3.7	Version Numbers and Ranges	61
3.7.1	Version Numbers	62
3.7.2	Versioning Bundles	63
3.7.3	Versioning Packages	63
3.7.4	Version Ranges	64
3.7.5	Versioning <code>Import-Package</code> and <code>Require-Bundle</code>	65
3.8	Class Loading in OSGi	66
3.9	JRE Packages	69
3.10	Execution Environments	70
3.11	Fragment Bundles	72
3.12	Class Space Consistency and “Uses” Constraints	73
4	Introduction to Services	75
4.1	Late Binding in Java	75
4.1.1	Dependency Injection Frameworks	76
4.1.2	Dynamic Services	77
4.2	Registering a Service	79
4.3	Unregistering a Service	81
4.4	Looking up a Service	84
4.5	Service Properties	86
4.6	Introduction to Service Trackers	88
4.7	Listening to Services	90
4.8	Tracking Services	92
4.9	Filtering on Properties	95
4.10	Cardinality and Selection Rules	96
4.10.1	Optional, Unary	98
4.10.2	Optional, Multiple	101
4.10.3	Mandatory, Unary	101
4.10.4	Mandatory, Multiple	101
5	Example: Mailbox Reader GUI	103

5.1	The Mailbox Table Model and Panel	103
5.2	The Mailbox Tracker	103
5.3	The Main Window	106
5.4	The Bundle Activator	109
5.5	Putting it Together	111
6	Concurrency and OSGi	115
6.1	The Price of Freedom	115
6.2	Shared Mutable State	117
6.3	Safe Publication	119
6.3.1	Safe Publication in Services	121
6.3.2	Safe Publication in Framework Callbacks	124
6.4	Don't Hold Locks when Calling Foreign Code	126
6.5	GUI Development	129
6.6	Using Executors	131
6.7	Interrupting Threads	138
6.8	Exercises	141
7	The Whiteboard Pattern and Event Admin	143
7.1	The Classic Observer Pattern	143
7.2	Problems with the Observer Pattern	144
7.3	Fixing the Observer Pattern	145
7.4	Using the Whiteboard Pattern	146
7.4.1	Registering the Listener	149
7.4.2	Sending Events	151
7.5	Event Admin	154
7.5.1	Sending Events	154
7.5.2	The Event Object	155
7.5.3	Receiving Events	158
7.5.4	Running the Example	159
7.5.5	Synchronous versus Asynchronous Delivery	161
7.5.6	Ordered Delivery	162
7.5.7	Reliable Delivery	162
7.6	Exercises	163
8	The Extender Model	165
8.1	Looking for Bundle Entries	166
8.2	Inspecting Headers	168
8.3	Tracking Bundles	169
8.4	Synchronous and Asynchronous Bundle Listeners	176
8.5	The Eclipse Extension Registry	181
8.6	Impersonating a Bundle	183
8.7	Conclusion	187
9	Configuration and Metadata	189

II Component Oriented Development	191
10 Component Oriented Development	193
11 Declarative Services	195
III Practical OSGi	197
12 Testing OSGi Bundles	199
13 Using Third-Party Libraries	201
14 Building Web Applications	203
IV Appendices	205
A Bundle Tracker	207
B ANT Build System for Bnd	211

List of Figures

1.1	The standard Java class loader hierarchy.	7
1.2	Example of an internal JAR dependency.	9
1.3	A broken internal JAR dependency.	10
1.4	Clashing Version Requirements	11
1.5	A typical J2EE class loader hierarchy.	14
1.6	The OSGi class loader graph.	17
1.7	Different Versions of the Same Bundle.	20
2.1	Adding Felix as a User Library in Eclipse	30
2.2	Creating a new Java project in Eclipse: adding the Felix library	32
2.3	A New OSGi Project, Ready to Start Work	33
2.4	Bundle Lifecycle	37
3.1	The runtime resolution of matching import and export.	57
3.2	The runtime resolution of a Required Bundle	60
3.3	Simplified OSGi Class Search Order	66
3.4	Full OSGi Search Order	68
4.1	Service Oriented Architecture	78
4.2	Updating a Service Registration in Response to Another Service	91
5.1	The Mailbox GUI (Windows XP and Mac OS X)	112
5.2	The Mailbox GUI with a Mailbox Selected	113
6.1	Framework Calls and Callbacks in OSGi	116
6.2	The Dining Philosophers Problem, Simplified	128
7.1	The Classic Observer Pattern	144
7.2	An Event Broker	146
7.3	A Listener Directory	147
7.4	The Event Admin Service	155
8.1	Inclusion Relationships of Bundle States	171
8.2	Bundle Transitions and Events	175
8.3	Synchronous Event Delivery when Starting a Bundle	179
8.4	Asynchronous Event Delivery after Starting a Bundle	180
8.5	Editing an Extension Point in Eclipse PDE	182

8.6 Editing an Extension in Eclipse PDE 183

B.1 OSGi Project Structure 212

List of Code Listings

2.1	A Typical OSGi MANIFEST.MF File	25
2.2	A Typical Bnd Descriptor File	27
2.3	Hello World Activator	35
2.4	Bnd Descriptor for the Hello World Activator	35
2.5	Bundle Counter Activator	42
2.6	Bnd Descriptor for the Bundle Counter	42
2.7	Heartbeat Activator	43
2.8	Hello Updater Activator	45
3.1	The Message Interface	49
3.2	The Mailbox Interface	50
3.3	Mailbox API Exceptions	51
3.4	Bnd Descriptor for the Mailbox API	52
3.5	String Message	54
3.6	Fixed Mailbox	55
3.7	MANIFEST.MF generated from <code>fixed_mailbox.bnd</code>	56
3.8	Bnd Sample: Controlling Bundle Contents	58
3.9	Bnd Sample: Controlling Imports	58
4.1	Naïve Solution to Instantiating an Interface	76
4.2	Welcome Mailbox Activator	79
4.3	Bnd Descriptor for the Welcome Mailbox Bundle	80
4.4	File Mailbox Activator	82
4.5	File Mailbox (Stub Implementation)	83
4.6	Bnd Descriptor for File Mailbox Bundle	83
4.7	Message Count Activator	85
4.8	Bnd Descriptor for the Message Counter Bundle	86
4.9	Adding Service Properties to the Welcome Mailbox	87
4.10	Message Count Activator — <code>ServiceTracker</code> version	88
4.11	Waiting for a Service	89
4.12	Database Mailbox Activator	93
4.13	Filter Building Utilities	97
4.14	Sample Usage of Filter Builder	97
4.15	Log Tracker	99
4.16	Sample Usage of Log Tracker	100
4.17	Multi Log Tracker	102

5.1	The Mailbox Table Model	104
5.2	Mailbox Panel	105
5.3	The Mailbox Tracker, Step One: Constructor	105
5.4	The Mailbox Tracker, Step Two: <code>addingService</code> method	107
5.5	The Mailbox Tracker, Step Three: <code>removedService</code> method	107
5.6	The Mailbox Reader Main Window	108
5.7	Conventional Java Approach to Launching a Swing Application	109
5.8	Using Bundle Lifecycle to Launch a Swing Application	110
5.9	Bnd Descriptor for the Mailbox Scanner	111
6.1	Thread Safety using Synchronized Blocks	118
6.2	Thread Safety using Read/Write Locks	120
6.3	Unsafe Publication	121
6.4	Dictionary Service interface	121
6.5	Unsafe Publication in a Service	121
6.6	Safe Publication in a Service	122
6.7	Connection Cache interface	123
6.8	Unsafe Connection Cache	123
6.9	Is This Bundle Activator Thread-safe?	124
6.10	Mailbox Registration Service Interface	127
6.11	Holding a lock while calling OSGi APIs	127
6.12	Avoiding holding a lock while calling OSGi APIs	130
6.13	Updating a Swing Control in Response to a Service Event	132
6.14	A Scala Utility to Execute a Closure in the Event Thread	133
6.15	Updating a Swing Control — Scala Version	133
6.16	Single-threaded execution	134
6.17	The <code>SerialExecutor</code> class	136
6.18	The Mailbox registration service using <code>SerialExecutor</code>	137
6.19	Registering a Thread Pool as an Executor Service	138
6.20	Server Activator	140
6.21	Exercise 2: Thread Pool Manager Interface	141
7.1	Mailbox Listener and Observable Mailbox Interfaces	144
7.2	Registering a Mailbox Listener	147
7.3	Mailbox Listener Tracker	148
7.4	Visitor Interface and Whiteboard Helper Class	149
7.5	Adding the <code>MailboxListener</code> Interface to <code>MailboxTableModel</code>	150
7.6	Mailbox Panel, with <code>MailboxListener</code> Registration	151
7.7	Growable Mailbox Activator and Timer Thread	152
7.8	Growable Mailbox	153
7.9	Random Price Generator	156
7.10	Event Admin Tracker	157
7.11	Bundle Activator for the Random Price Generator	157
7.12	Stock Ticker Activator	160
7.13	Bnd Descriptors for the Random Price Feed and Stock Ticker	160

8.1	Help Provider Service Interface	165
8.2	Scanning a Bundle for Help Documents	166
8.3	A Help Index File, <code>index.properties</code>	167
8.4	Scanning a Bundle for Help Documents with Titles (1)	167
8.5	The <code>Pair</code> Class	168
8.6	<code>MANIFEST.MF</code> for a Help Provider	169
8.7	Scanning a Bundle for Help Documents with Titles (2)	170
8.8	The <code>HelpExtender</code> Class	173
8.9	Testing Multiple Bundle States in One <code>if</code> Statement	175
8.10	Shell Command for Testing the Help Extender	177
8.11	Activator for the Help Extender Bundle	178
8.12	Bnd Descriptor for the Help Extender Bundle	178
8.13	Bnd Descriptor for a Sample Help Provider	178
8.14	An Eclipse <code>plugin.xml</code> File	182
8.15	Mailbox Service Extender	185
8.16	Activator for the Mailbox Service Extender	186
8.17	Bnd Descriptor for the Mailbox Service Extender	186
8.18	Minimal Mailbox Class	187
8.19	Bnd Descriptor for a “Declarative” Mailbox Service Bundle	187
A.1	<code>BundleTracker</code>	207
B.1	<code>build.properties</code>	211
B.2	<code>build.xml</code>	213

Preface

Version


This book was generated on January 11, 2009 from the following tag:

- Preview-20080922-0200

Please include the tag when reporting errors or amendments.

License

This book is licensed under a Creative Commons Attribution Non-commercial Share Alike License.



You are free:

- to copy, distribute, display, and perform the work
- to make derivative works

Under the following conditions:

- *Attribution.* You must give the original author credit.
- *Non-Commercial.* You may not use this work for commercial purposes.
- *Share Alike.* If you alter, transform, or build upon this work, you may distribute the resulting work only under a licence identical to this one.

For any reuse or distribution, you must make clear to others the licence terms of this work. Any of these conditions can be waived if you get permission from the copyright holder. Nothing in this license impairs or restricts the author's moral rights.

Your fair dealing and other rights are in no way affected by the above. This is a human-readable summary of the [Legal Code \(the full licence\)](#)¹.

¹<http://creativecommons.org/licenses/by-nc-sa/2.0/uk/legalcode>

Part I

Nuts and Bolts

1 Introduction

Consider the following question. In any large engineering project, for example the design of a new bridge, skyscraper, or jet airliner, what is nearly always the most difficult challenge to overcome?

The answer is *Complexity*.

The Boeing 747-400 “Jumbo Jet” has six million parts^[1], of which half are fasteners. It contains 171 miles (274 km) of wiring and 5 miles (8 km) of tubing. Seventy-five thousand engineering drawings were used to produce the original 747 design. It is a very, *very* complex machine, and because of this complexity, no single person can have a complete understand of how a Boeing 747 works. Yet, this is a machine first designed in the 1960s — modern aircraft have multiplied the level of complexity many times over. The only approach that will allow human beings to design such incredible machines is to break them down into smaller, more understandable modules.

Modularity enables several important benefits:

Division of Labour. We can assign separate individuals or groups to work on separate modules. The people working on a module will have a thorough understanding of their own module but not all the others. For example, the designer of the entertainment system does not need to know anything about how the landing gear works (and *vice versa!*) but can concentrate all her efforts on building the best possible entertainment system.

Abstraction. We can now think about the aircraft as an abstract whole, without needing a complete understanding of every part. For example, we grasp that a 747 is able to fly due to the lift provided by the wings and the forward thrust of the engines, even if we do not understand exactly how fuel is supplied to the engines or how many motors are required to move the flaps up and down.

Reuse. Given the amount of effort that goes into designing even some of the smaller components of an aircraft, it would be a shame to have to start from scratch when we need an identical or very similar component in another aircraft, or even in another part of the same aircraft. It would be helpful if we could reuse components with minimal alterations.

Ease of Maintenance and Repair. It would be a shame to have to scrap an entire aeroplane over a single burst tyre or torn seat cover. Modular

designs allow for failed modules to be removed and either repaired or replaced without affecting the rest of the machine.

It's debatable whether the production of software can be characterised as an engineering discipline, but nevertheless the complexity of software rivals what is seen in other fields. Modern software is incredibly complex, and furthermore is accelerating in its complexity. For example, the onboard computers of the NASA Space Shuttle contained half a million lines of code, but today the DVD player in your living room contains around one million lines of code. Microsoft Windows XP is estimated to contain 40 million lines, and Windows Vista 50 million. A BMW 7-series car can contain up to 50 networked computers.

Just like aircraft engineers, software professionals are in the business of creating large and complex machines, which we can only hope to understand by breaking them into modules.

The Java™ programming language is one of the most popular languages today for building both large, enterprise applications and also small but widely deployed mobile applications. However Java alone does not support modularity in any useful way. . . we will see in the next section why Java's existing mechanisms fail to deliver all four of the above listed benefits of modularity. However, Java's great strength is its flexibility, which has allowed a powerful module system to be built on top.

That module system is called OSGi. OSGi is nothing more nor less than the way to build modular applications in Java.

1.1 What is a Module?

So, what should a software module look like? It should have the following properties:

Self-Contained. A module is a logical whole: it can be moved around, installed and uninstalled as a single unit. It is not an atom — it consists of smaller parts — but those parts cannot stand alone, and the module may cease to function if any single part is removed.

Highly Cohesive. Cohesion is a measure of how strongly related or focussed the responsibilities of a module are. A module should not do many unrelated things, but stick to one logical purpose and fulfil that purpose well.

Loosely Coupled. A module should not be concerned with the internal implementation of other modules that it interacts with. Loose coupling allows us to change the implementation of one module without needing to update all other modules that use it (along with any modules that use *those* modules, and so on).

To support all three properties, it is vital for modules to have a *well-defined interface* for interaction with other modules. A stable interface enforces logical boundaries between modules and prevents access to internal implementation details. Ideally, the interface should be defined in terms of what each module offers to other modules, and what each module requires from other modules.

1.2 The Problem(s) with JARs

The standard unit of deployment in Java is the JAR file, as documented by the JAR File Specification[2]. JARs are archive files based on the ZIP file format, allowing many files to be aggregated into a single file. Typically the files contained in the archive are a mixture of compiled Java class files and resource files such as images and documents. Additionally the specification defines a standard location within a JAR archive for metadata — the `META-INF` folder — and several standard file names and formats within that directly, most important of which is the `MANIFEST.MF` file.

JAR files typically provide either a single library or a portion of the functionality of an application. Rarely is an entire application delivered as a single JAR, unless it is very small. Therefore, constructing Java applications requires composing together many JAR files: large applications can have a JAR count in double or even triple figures. And yet the Java Development Kit (JDK) provides only very rudimentary tools and technologies for managing and composing JARs. In fact the tools available are so simplistic that the term “JAR Hell”¹ has been coined for the problem of managing JARs, and the strange error messages that one encounters when things go wrong.

The three biggest problems with JAR files as a unit of deployment are as follows:

- There is no runtime concept that corresponds to a JAR; they are only meaningful at build-time and deploy-time. Once the Java Virtual Machine is running, the contents of all the JARs are simply concatenated and treated as a single, global list: the so-called “Classpath”. This model scales very poorly.
- They contain no standard metadata to indicate their dependencies.
- They are not versioned, and multiple versions of JARs cannot be loaded simultaneously.
- There is no mechanism for information hiding between JARs.

¹A reference to the term **DLL Hell**: a comparable problem with managing DLL files on the Microsoft Windows operating system.

1.2.1 Class Loading and the Global Classpath

The term “classpath” originates from the command-line parameter that is passed to the `java` command when running simple Java applications from the command shell (or the DOS prompt in Windows terms), or more usually from some kind of launcher script. It specifies a list of JAR files and directories containing compiled Java class files. For example the following command launches a Java application with both `log4j.jar` and the `classes` directory on the classpath. First as it would be under UNIX (including Mac OS X):

```
java -classpath log4j.jar:classes org.example.HelloWorld
```

And then the DOS/Windows version:

```
java -classpath log4j.jar;classes org.example.HelloWorld
```

The final parameter is the name of the “main” class to execute, which we will assume has been compiled into the `org/example/HelloWorld.class` file in the `classes` directory. Somehow the Java Virtual Machine (JVM) must load the bytes in that class file and transform them into a `Class` object, on which it can then execute the static `main` method. Let’s look at how this works in a standard Java runtime environment (JRE).

The class in Java that loads classes is `java.lang.ClassLoader`, and it has two responsibilities:

1. Finding classes, i.e. the physical bytes on disk, given their logical class names.
2. Transforming those physical bytes into a `Class` object in memory.

We can extend the `java.lang.ClassLoader` class and provide our own implementation of the first part, which allows us to extend the JRE to load code from a network or other non-file-based storage system. However the second part, i.e. transforming physical class bytes into `Class` objects, is implemented in `java.lang.ClassLoader` by the `defineClass` method, which is both `native` and `final`. In other words, this functionality is built into the JRE and cannot be overridden.

When we run the command above, the JRE determines that it needs to load the class `org.example.HelloWorld`. Because it is the “main” class, it consults a special `ClassLoader` named the application class loader. The first thing the application class loader does is ask its “parent” to load the class.

This illustrates a key feature of Java class loading called parent-first delegation. Class loaders are organised into a hierarchy, and by default all class loaders

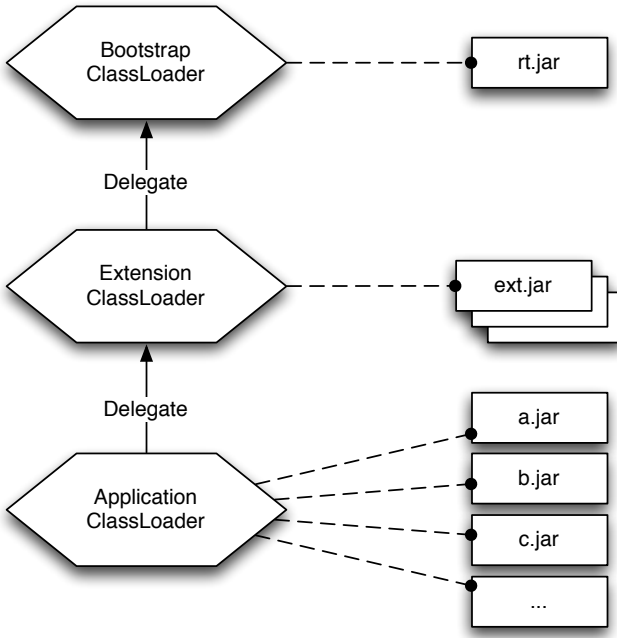


Figure 1.1: The standard Java class loader hierarchy.

first ask their parent to load a class; only when the parent responds that it knows nothing about the specified class does the class loader attempt to find the class itself. Furthermore the parent will also delegate to its own parent, and so on until the top of the hierarchy is reached. Therefore classes will always be loaded at the highest possible level of the hierarchy.

Figure 1.1 shows the three standard class loaders found in a conventional Java application. The bootstrap class loader sits at the top of the tree, and it is responsible for loading all the classes in the base JRE library, for example everything with a package name beginning with `java`, `javax`, etc. Next down is the extension class loader, which loads from “extension” libraries, i.e. JARs which are not part of the base JRE library but have been installed into the `libext` directory of the JRE by an administrator. Finally there is the aforementioned application class loader, which loads from the “classpath”.

Returning to our “HelloWorld” example, we can assume that the `HelloWorld` class is not present in the JRE base or extension libraries, and therefore it will not be found by either the bootstrap or extension class loaders. Therefore the application class loader will try to find it, and it does that by looking inside each entry in the classpath in turn and stopping when it finds the first match. If a match is not found in any entry, then we will see the familiar

ClassNotFoundException.

The `HelloWorld` class is probably not inside `log4j.jar`, but it will be found in the `classes` directory, and loaded by the class loader. This will inevitably trigger the loading of several classes that `HelloWorld` depends on, such as its super-class (even if that is only `java.lang.Object`), any classes used in its method bodies and so on. For each of those classes, the whole procedure described above is repeated.

To recap:

1. The JRE asks the application class loader to load a class.
2. The application class loader asks the extension class loader to load the class.
3. The extension class loader asks the bootstrap class loader to load the class.
4. The bootstrap class loader fails to find the class, so the extension class loader tries to find it.
5. The extension class loader files to find the class, so the application class loader tries to find it, looking first in `log4j.jar`.
6. The class is not in `log4j.jar` so the class loader looks in the `classes` directory.
7. The class is found and loaded, which may trigger the loading of further classes — for each of these we go back to step 1.

1.2.2 Conflicting Classes

The process for loading classes in Java does at least work for much of the time. However, consider what would happen if we accidentally added a JAR file to our classpath containing an older version of `HelloWorld`. Let's call that file `obsolete.jar`.

```
java -classpath obsolete.jar:log4j.jar:classes \
    org.example>HelloWorld
```

Since `obsolete.jar` appears before `classes` in the classpath, and since the application class loader stops as soon as it finds a match, this command will always result in the old version of `HelloWorld` being used. The version in our `classes` directory will *never* be used. This can be one of the most frustrating problems to diagnose in Java: it can appear that a class is not doing what it should be doing, and the changes we are making to it are not having any effect, simply because that class is being shadowed by another with the same name.




Figure 1.2: Example of an internal JAR dependency.

Perhaps this scenario does not seem very likely, and it's true that in a trivial application like this, such a mistake is easily avoided. However, as we discussed, a large Java application can consist of tens or even hundreds of individual JAR files, all of which must be listed on the classpath. Also, JAR files frequently have obscure names which give little clue as to their contents. Faced with such problems it becomes much more likely — perhaps even inevitable — that the classpath will contain conflicting names. For example, one of the commonest cause of conflicts is including two different versions of the same library.

By simply searching classpath entries in order and stopping on the first match, the JRE reduces JARs to mere lists of class files, which are dumped into a single flat list. The boundaries of the “inner” lists — that is, the JARs themselves — are essentially forgotten. For example, suppose a JAR contains two internal classes, `Foo` and `Bar`, such that `Foo` uses `Bar`. This is shown in Figure 1.2.

Presumably, `Foo` expects to resolve to the version of `Bar` that is shipped alongside it. It has almost certainly been built and tested under that assumption. However at runtime the connection between the two classes is arbitrarily broken because another JAR happens to contain a class called `Bar`, as shown in Figure 1.3

1.2.3 Lack of Explicit Dependencies

Some JAR files are “standalone”, i.e. they provide functionality without depending on any other libraries save for the base JRE libraries. However, many build on the functionality offered by other JARs, and therefore they can only be used if they are deployed alongside those other JARs. For example, the Apache Jakarta Commons HttpClient^[3] library depends on both Commons Codec and Commons Logging, so it will not work without `commons-logging.jar` and `commons-codec.jar` being present on our classpath.

But how do we know that such a dependency exists? Fortunately `HttpClient` is well documented, and the dependency is explicitly noted on the project web




Figure 1.3: A broken internal JAR dependency.

site. But many libraries do not have such good documentation. Instead the dependency is implicit: lurking inside the class file in the JAR, ready to cause a `ClassNotFoundException` when we try to use it.

In fact, even the good documentation exemplified by `HttpClient` is not really good enough. What we want is a standard way to declare dependencies, preferably right in the JAR file, such that the declarations can be analysed easily by tools.

There was an early attempt in Java to supply such information through the `Class-Path` attribute, which can be specified in `MANIFEST.MF`. Unfortunately this mechanism is almost entirely useless, because it only allows one to list further JAR files to be added to the classpath using absolute file-system paths, or paths relative to the file-system location of the JAR in question.

1.2.4 Lack of Version Information

The world does not stand still, and neither do the libraries available to us as Java programmers. New libraries and new versions of existing libraries appear all the time.

Therefore it is not enough merely to indicate a dependency on a particular library: we must also know which *version* is required. For example, suppose we determine somehow that a JAR has a dependency on `Log4J`. Which version of `Log4J` do we need to supply to make the JAR work? The download site lists 25 different versions (at time of writing). We cannot simply assume the latest version, since the JAR in question may not have been tested against it,

and in the case of Log4J there is an experimental version 1.3.x that almost nobody uses.

Again, documentation sometimes comes to the rescue (e.g., “this library requires version 1.2.10 of Log4J or greater”), but unfortunately this kind of information is very rare, and even when it does exist it is not in a format that can be used by tools. So we need a way to tag our explicitly declared dependencies with a version. . . . in fact, we need to specify a version *range* because depending on a single specific version of a library would make our system brittle.

Versions cause other problems. Suppose our application requires two libraries, *A* and *B*, and both of these libraries depend in turn upon a third library, *C*. But they require different versions: library *A* requires version 1.2 or greater of *C*, but library *B* requires version 1.1 of *C*, and will not work with version 1.2! Figure 1.4 illustrates the problem.

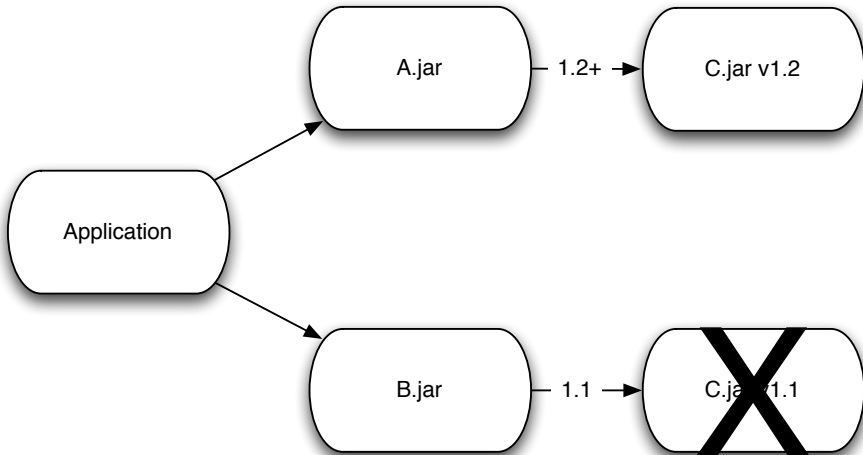


Figure 1.4: Clashing Version Requirements

This kind of problem simply cannot be solved in traditional Java without rewriting either *A* or *B* so that they both work with the same version of *C*. The reason for this is simply the flat, global classpath: if we try to put both versions of *C* on the classpath, then only the first version listed will be used, so both *A* and *B* will have to use it. Note that version 1.2 in this example does *not* win because it is the higher version; it wins merely because it was placed first in the classpath.

However, some classes from version 1.1 may still be visible if they have different names from the classes in 1.2! If *B* manages to use one of those classes, then

that class will probably attempt to use other classes that are shadowed by version 1.2, and it will get the 1.2 version of those classes rather than the 1.1 version that it has been compiled against. This is an example of the class loading conflict from Figure 1.2. The result is likely to be an error such as `LinkageError`, which few Java developers know how to deal with.

1.2.5 Lack of Information Hiding Across JARs

All Object Oriented Programming languages offer various ways of hiding information across class and module boundaries. Hiding — or *encapsulating* — internal details of a class is essential in order to allow those internal details to change without affecting clients of the class. As such, Java provides four access levels for members of a class, which include both fields and methods:

- `public` members are visible to everybody.
- `protected` members are visible to subclasses and other classes in the same package.
- `private` members are visible only within the same class.
- Members not declared with any of the three previous access levels take the so-called *default* access level. They are visible to other classes within the same package, but not classes outside the package.

For classes themselves, only the `public` and default access levels are available. A public class is visible to every class in every other package; a default access class is only available to other classes within the same package.

There is something missing here. The above access modifiers relate to visibility across packages, but the unit of deployment in Java is not a package, it is a JAR file. Most JAR files offering non-trivial APIs contain more than one package (`HttpClient` has eight), and generally the classes within the JAR need to have access to the classes in other packages of the same JAR. Unfortunately that means we must make most of our classes `public`, because that is the only access modifier which makes classes visible across package boundaries.

As a consequence, all those classes declared `public` are accessible to clients outside the JAR as well. Therefore the whole JAR is effectively public API, even the parts that we would prefer to keep hidden. This is another symptom of the lack of any runtime representation for JAR files.

1.2.6 Recap: JARs Are Not Modules

We've now looked in detail at some specific problems with JAR files. Now let's recap why they do not meet the requirements for a module system.

Simply, JAR files have almost none of the characteristics of a module as described in Section 1.1. Yes, they are a unit that can be physically moved around. . . but having moved a JAR we have no idea whether it will still work, because we do not know what dependencies might be missing or incorrect.

JARs are often tightly coupled: thanks to the lack of information hiding, clients can easily use internal implementation details of a JAR, and then break when those details change. And those are just the clients who break encapsulation intentionally; other clients may accidentally use the internals of a JAR thanks to a classpath conflict. Finally, JARs often have low cohesion: since a JAR does not really exist at runtime, we might as well throw any functionality into whichever JAR we like.

Of course, none of this implies that building modular systems in Java is impossible². It simply means that Java provides no assistance towards the goal of modularity. Therefore building modular systems is a matter of *process* and *discipline*. Sadly, it is rare to see large Java applications that are modular, because a disciplined modular approach is usually the first thing discarded when an important deadline looms. Therefore most such applications are a mountain of “spaghetti code”, and a maintenance nightmare.

1.3 J2EE Class Loading

The Java 2 Enterprise Edition (J2EE) specification defines a platform for distributed, multi-tier computing. The central feature of the J2EE architecture is the so-called “application server” which hosts multiple application components and offers enterprise-class services to them such as transaction management, security and high availability. Application servers are also required to have a deployment system so that applications can be deployed and un-deployed without restarting the server and without interfering with each other.

These requirements meant that the simplistic class loading hierarchy of Figure 1.1, as used by standalone Java applications, was not sufficient. With a single flat classpath, the classes from one application could easily interfere with other applications. Therefore J2EE application servers use a more complex tree, with a branch for each deployed application. The precise layout of the tree depends on the individual application server, since it is not mandated by the specification, but Figure 1.5 shows the approximate layout used by most servers.

J2EE applications are deployed as “Enterprise ARchive” (EAR) files, which are ZIP files containing a metadata file — `application.xml` — plus one or more of each of the following kind of file:

²Although certain scenarios are indeed impossible, such as the side-by-side usage of different versions of the same library.




Figure 1.5: A typical J2EE class loader hierarchy.

- Plain Java library JAR files
- JAR files containing an EJB application (EJB-JARs)
- “Web ARchive” (WAR) files, containing classes implementing Web functionality, such as Servlets and JSPs.

The plain Java library JAR files contain classes that are supposed to be available to all of the EJBs and Web artifact within the EAR. They are therefore loaded by the EAR Class Loader, at the top of the sub-tree for the deployed application.

Referring back to the class loading procedure described in section 1.2.1, it should be clear that, in a branching tree, an individual class loader can only load classes defined by itself or its ancestors; it cannot load classes sideways in the tree, i.e. from its siblings or cousins. Therefore classes that need to be shared across both the EJBs and Web artifacts must be pushed up the tree, into the EAR Class Loader. And if there are classes we wish to share across multiple deployed applications, we must push them up into the system application class loader. Usually this is done by configuring the way the application server itself is started, adding JARs to the global classpath.

Unfortunately, libraries pushed up to that level can no longer be deployed and un-deployed at will. Also, they become available to all deployed applications, even the ones that do not want or need to use them. They cause class conflicts: classes found higher in the tree always take precedence over the classes that are shipped with the application. In fact, every application in the server must now use the *same version* of the library.

Because of these problems, J2EE developers tend to avoid sharing classes across multiple applications. When a library is needed by several applications, it is simply shipped multiple times as part of the EAR file for each one. The end result is a “silo” or “stovepipe” architecture: applications within the J2EE server are little more than standalone systems, completely vertically integrated from the client tier to the database, but unable to horizontally integrate with each other. This hinders collaboration amongst different business areas.

1.4 OSGi: A Simple Idea

OSGi is the module system for Java. It defines a way to create true modules and a way for those modules to interact at runtime.

The central idea of OSGi is in fact rather simple. The source of so many problems in traditional Java is the global, flat classpath, so OSGi takes a completely different approach: each module has its own classpath, separate from the classpath of all other modules. This immediately eliminates almost all of problems that were discussed, but of course we cannot simply stop there:

modules do still need to work together, which means sharing classes. The key is *control* over that sharing. OSGi has very specific and well-defined rules about how classes can be shared across modules, using a mechanism of explicit imports and exports.

So, what does an OSGi module look like? First, we don't call it a module: in OSGi, we refer to *bundles*.

In fact a bundle is just a JAR file! We do not need to define a new standard for packaging together classes and resources: the JAR file standard works just fine for that. We just need to add a little metadata to promote a JAR file into a bundle. The metadata consists of:

- The *name* of the bundle. We provide a “symbolic” name which is used by OSGi to determine the bundle's unique identity, and optionally we can also provide a human-readable, descriptive name.
- The *version* of the bundle.
- The list of *imports* and *exports*. We will see shortly exactly what is being imported and exported.
- Optionally, information about the minimum Java version that the bundle needs to run on.
- Miscellaneous human-readable information such as the vendor of the bundle, copyright statement, contact address, etc.

These metadata are placed inside the JAR file in a special file called `MANIFEST.MF`, which is part of all standard JAR files and is meant for exactly this purpose: it is a standard place to add arbitrary metadata.

The great advantage of using standard JAR files as OSGi bundles is that a bundle can also be used everywhere that a JAR file is expected. Not everybody uses OSGi (yet...) but since bundles are just JAR files, they can be used outside of the OSGi runtime. The extra metadata inside `MANIFEST.MF` is simply ignored by any application that does not understand it.

1.4.1 From Trees to Graphs

What does it mean to provide a separate classpath for each bundle? Simply, we provide a class loader for each bundle, and that class loader can see the classes and resources inside the bundle's JAR file. However in order for bundles to work together, there must be some way for class loading requests to be delegated from one bundle's class loader to another.

Recall that in both standard Java and Enterprise Edition, class loaders are arranged in a hierarchical tree, and class loading requests are always delegated upwards, to the parent of each class loader. Recall also that this arrangement




Figure 1.6: The OSGi class loader graph.

does not allow for sharing of classes horizontally across the tree i.e., between siblings or cousins. To make a library available to multiple branches of the tree it must be pushed up into the common ancestor of those branches, but as soon as we do this we force everybody to use that version of that library, whether they like it or not.

Trees are simply the wrong shape: what we really need is a *graph*. The dependency relationship between two modules is not a hierarchical one: there is no “parent” or “child”, only a network of providers and users. Class loading requests are delegated from one bundle’s class loader to another’s based on the dependency relationship between the bundles.

An example is shown in Figure 1.6. Here we have five libraries with a complex set of interdependencies, and it should be clear that we cannot easily force this into a hierarchical tree shape.

The links between bundles are based on imported and exported packages. That is, Java packages such as `javax.swing` or `org.apache.log4j`.

Suppose bundle *B* in Figure 1.6 contains a Java package named `org.foo`. It may choose to export that package by declaring it in the exports section of its `MANIFEST.MF`. Bundle *A* may then choose to import `org.foo` by declaring it in the imports section of its `MANIFEST.MF`. Now, the OSGi framework will take responsibility for matching up the import with a matching export: this is known as the resolution process. Resolution is quite complex, but it is implemented by the OSGi framework, not by bundles themselves. All that we need to do, as developers of bundles, is write a few simple declarative

statements.

Once an import is matched up with an export, the bundles involved are “wired” together for that specific package name. What this means is that when a class load request occurs in bundle *A* for any class in the `org.foo` package, that request will immediately be delegated to the class loader of bundle *B*. Other imports in *A* may be wired to other bundles, for example *A* may also import the package `org.bar` that is exported by bundle *C*, so any loading requests for classes in the `org.bar` package will be delegated to *C*’s class loader. This is extremely efficient: whereas in standard Java class load events invariably involve searching through a long list of classes, OSGi class loaders generally know immediately where to find a class, with little or no searching.

What happens when resolution fails? For example, what if we forget to include bundle *B* with our application, and no other bundle offers package `org.foo` to satisfy the import statement in bundle *A*? In that case, *A* will not resolve, and cannot be used. We will also get a helpful error message telling us exactly why *A* could not be resolved. Assuming our bundles are correctly constructed (i.e., their metadata is accurate) then we should never see errors like `ClassNotFoundException` or `NoClassDefFoundError`³. In standard Java these errors can pop up at any time during the execution of an application, for example when a particular code path is followed for the first time. By contrast an OSGi-based application can tell us at start-up that something is wrong. In fact, using simple tools to inspect the bundle metadata, we can know about resolution errors in a set of bundles *before* we ever execute the application.

1.4.2 Information Hiding in OSGi Bundles

Note that we always talk about matching up an import with an export. But why are the export declarations even necessary? It should be possible simply to look at the contents of a bundle JAR to find out what packages are contained within it, so why duplicate this information in the exports section of the `MANIFEST.MF`?

The answer is, we don’t necessarily want to export all packages from a bundle for reasons of information hiding as discussed in Section 1.2.5. In OSGi, only packages that are explicitly exported from a bundle can be imported and used in another bundle. Therefore all packages are “bundle private” by default, making it easy for us to hide the internal implementation details of our bundles from clients.

³An exception to this is where dynamic reflection is used to load arbitrary classes at runtime, a topic which is discussed in Chapter ??

1.4.3 Versioning and Side-by-Side Versions

OSGi does not merely offer dependencies based on package names: it also adds versioning of packages. This allows us to cope with changes in the released versions of libraries that we use.

Exports of packages are declared with a version attribute, but imports declare a version *range*. This allows us to have a bundle depend on, e.g., version 1.2.0 through to version 2.1.0 of a library. If no bundle is exporting a version of that package that falls within the specified range, then our bundle will fail to resolve, and again we get a helpful error message telling us what is wrong.

We can even have different versions of the same library side-by-side in the same application, so the scenario described in Section 1.2.4 (wherein two libraries each need to use different versions of a third library) will work under OSGi.

Note that, because exports are declared on each exported package, there is no need for all exports from a bundle to be the same version. That is, a single bundle can export both version 1.2 of `org.foo` and version 3.5 of `org.bar`. Of course, it cannot export two versions of the *same* package.

1.5 Dynamic Modules

OSGi is not just the module system for Java. It is the *dynamic* module system for Java. Bundles in OSGi can be installed, updated and uninstalled without taking down the entire application. This makes it possible to, for example, upgrade parts of a server application — either to include new features or to fix bugs found in production — without affecting the running of other parts. Also, desktop applications can download new versions without requiring a restart, so the user doesn't even need to be interrupted.

To support dynamic modules, OSGi has a fully developed lifecycle layer, along with a programming model that allows “services” to be dynamically published and consumed across bundle boundaries. We will look at these facilities in depth in later chapters.

Incidentally, many developers and system administrators are nervous or sceptical about OSGi's dynamic capabilities. This is perfectly understandable after the experience of J2EE, which offered limited and unreliable hot deployment of EAR modules. These people should still consider using OSGi anyway for its great modularity benefits, and feel free to ignore the lifecycle layer. Nevertheless, OSGi's dynamic capabilities are not mere hype: they really do work, and some experimentation will confirm that.




Figure 1.7: Different Versions of the Same Bundle.

1.6 The OSGi Alliance and Standards

OSGi is a standard defined by an Alliance of around forty companies. Its specifications are freely available and are comprehensive enough that a highly compliant implementation can be written using only the documents for reference.

Two of the commonest questions about the name OSGi are: what does it stand for, and why is the “i” lower-case? Here is the definitive answer to both questions: officially OSGi *doesn't stand for anything!* However, it *used to* stand for “Open Service Gateway initiative”, and this is the source of the lower-case letter “i”, since “initiative” was not considered to be stylistically part of the brand name. But the long name is now deprecated, since OSGi has expanded far beyond its original role in home gateways. As a result, the OSGi name is rather odd, but it has the great advantage of being easily Google-able since it seems not to be a word (or even part of a word) in any language. In speech, the name is always spelt out (“Oh Ess Gee Eye”) rather than pronounced as a word (e.g., “Ozjee”).

The OSGi Alliance’s role is to define the specification for new releases of the platform, and to certify implementations of the current release of the specification. The technical work is done by a number of Expert Groups (EGs) which include the Core Platform Expert Group (CPEG), Mobile (MEG), Vehicle (VEG) and Enterprise (EEG) Expert Groups. In this book we will mostly look at the work of the CPEG.

1.7 OSGi Implementations

Several independently implemented OSGi frameworks exist today, including four that are available as open source software.

Equinox [4] is the most widely deployed OSGi framework today owing to its use in the core runtime of Eclipse. It can also be found in many in-house custom applications as well as packaged products such as Lotus Notes and IBM WebSphere Application Server. Equinox implements Release 4.1 of the OSGi specifications and is licensed under the Eclipse Public License (EPL)[5].

Knopflerfish [6] is a popular and mature implementation of both OSGi Release 3 and Release 4.1 specifications. It is developed and maintained by Makewave AB, a company based in Sweden, and is licensed under a BSD-style license. Makewave also offers Knopflerfish Pro, a commercially supported version of Knopflerfish.

Felix [7] is a community implementation of the OSGi Release 4.x under the Apache Group umbrella. It is designed particularly for compactness and ease of embedding, and is the smallest (in terms of minimal JAR size) of the Release 4 implementations. It is licensed under the Apache License Version 2.0.

Concierge [8] is a very compact and highly optimized implementation of OSGi Release 3. This makes it particularly suited to resource-constrained platforms such as mobile phones. Concierge is licensed under a BSD-style license. However, OSGi Release 3 is not covered by this book⁴.

All of the example code and applications in this book should work on any of the three listed Release 4 implementations, except where explicitly noted. However, the way we work with each of these frameworks — for example the command line parameters, the built-in shell, the management of bundles — differs enough that it would be awkward to give full instructions for all three. Therefore it is necessary to pick a single implementation for pedagogical purposes, and for reasons explained later we will work with Felix.

1.8 Alternatives to OSGi

At its core, OSGi is very simple and, as with all good and simple ideas, many people have independently invented their own versions at different times and places. None of these have achieved the maturity or widespread usage that

⁴Also note that Concierge has not been certified compliant with OSGi R3, since OSGi Alliance rules only allow implementations of the *current* specification release to be certified.

OSGi now enjoys, but it is informative to review some of these alternatives. At the very least we may learn something about why OSGi has made the design choices it has made, but we may also find good ideas that can be brought into OSGi.

1.8.1 Build Tools: Maven and Ivy

Maven and Ivy are both popular tools that have some characteristics of a module system, but they are build-time tools rather than runtime frameworks. Thus they do not compete directly with OSGi, in fact they are complementary and many developers are using Maven or Ivy to build OSGi-based systems.

Maven is a complete, standalone build tool, whereas Ivy is a component that can be integrated into an ANT-based build. Both tools attempt to make JARs more manageable by adding modular features to them. Principally this means dependencies: both allow us to specify the versioned dependencies of a JAR using metadata in XML files. They use this information to download the correct set of JARs and construct a compile-time classpath.

However, as they do not have any runtime features neither Maven nor Ivy can solve the runtime problems with JARs, such as the flat global classpath, the lack of information hiding, and so on. Also the metadata formats used by these tools is unfortunately not compatible with the format used by OSGi, so if we use Maven or Ivy to build an OSGi-based system we typically have to specify the metadata twice: once for OSGi and once for the build tool. However, some efforts are currently being made to better integrate Maven with OSGi.

1.8.2 Eclipse Plug-in System

As already noted, the Eclipse IDE and platform are based on an implementation of OSGi. However this was not always the case: prior to version 3.0, Eclipse used its own custom module system.

In Eclipse terminology, a module is a “plug-in”. In fact, Eclipse developers often still use the term plug-in as an alternative name for an OSGi bundle. In the old Eclipse system, a plug-in was a directory containing a file at the top level named `plugin.xml`. This file contained metadata that was broadly similar to the metadata in an OSGi manifest: the name of the plug-in, vendor, version, exported packages and required plug-ins.

Notice a key difference here. In the Eclipse plug-in system, dependencies were not declared at the level of Java packages but of whole plug-ins. We would declare a dependency on a plug-in based on its ID, and this would give us access to *all* of the exported packages in that plug-in. OSGi actually supports

whole-bundle dependencies also, but the use of this capability is frowned upon for reasons we will examine in Chapter 3.

The biggest deficiency of the Eclipse plug-in system was its inability to install, update or uninstall plug-ins dynamically: whenever the plug-in graph changed, a full restart was required. In early 2004 the core Eclipse developers began a project, code-named Equinox, to support dynamic plug-ins. They intended to do this either by enhancing the existing system or selecting an existing module system and adapting it to Eclipse. In the end, OSGi was selected and Equinox became a leading implementation of it.

1.8.3 JSR 277

Java Specification Request (JSR) number 277 is titled *Java Module System*^[9] and as such one would expect it to attempt to solve a similar set of problems to OSGi. However JSR 277 tackles them in a different way to OSGi.

The most important point to note is that, at the time of writing, JSR 277 is an incomplete specification with no implementation yet. It is scheduled to be included with Java 7 but it is unclear when (or even if!) Java 7 will be released. An Early Draft Review (EDR) of JSR 277 was released in October 2006, which is the best information currently available.

Like the old Eclipse plug-in system, JSR 277 does not support package-level dependencies, instead using whole-module dependencies. However JSR 277 differs further by using programmatic resolution “scripts” rather than declarative statements to resolve module dependencies. This allows for extreme flexibility, but it’s debatable whether such flexibility is ever necessary or desirable. Programmatic code can return different results at different times, so for example we could write a module that resolves successfully only on Tuesday afternoons! Therefore we completely lose the ability to use static dependency analysis tools.

Furthermore, JSR 277 is not dynamic, it requires the Java Virtual Machine to be restarted in order to install, update or uninstall a module.

In a sense, JSR 277 is an affirmation from its sponsors (principally Sun) that modularity is important and currently missing from standard Java. Sadly, JSR 277 comes many years later than OSGi, includes some questionable design features, and is substantially less ambitious despite its opportunity to change the underlying Java runtime. Therefore we hope that JSR 277 will at the very least be compatible with OSGi, since it currently appears to represent a significant step backwards for modularity in Java.

2 First Steps in OSGi

OSGi is a module system, but in OSGi we refer to modules as “bundles”. In this chapter we will look at the structure of a bundle, how it depends on other bundles, and how to create an OSGi project using standard Java tools.

2.1 Bundle Construction

As discussed in the Introduction, an OSGi bundle is simply a JAR file, and the only difference between a bundle JAR and a “plain” JAR is a small amount of metadata added by OSGi in the form of additional headers in the `META-INF/MANIFEST.MF` file. Bundles can be used outside of an OSGi framework, for example in a plain Java application, because applications are required to ignore any attributes in the manifest that they do not understand.

Listing 2.1 A Typical OSGi MANIFEST.MF File

```
1 Manifest-Version: 1.0
2 Created-By: 1.4.2_06-b03 (Sun Microsystems Inc.)
3 Bundle-ManifestVersion: 2
4 Bundle-Name: My First OSGi Bundle
5 Bundle-SymbolicName: org.osgi.example1
6 Bundle-Version: 1.0.0
7 Bundle-RequiredExecutionEnvironment: J2SE-1.5
8 Import-Package: javax.swing
```

Listing 2.1 shows an example of a manifest containing some of the most common attributes used in OSGi. The lines in *italic font* are standard in all JAR file manifests, although only the `Manifest-Version` attribute is mandatory and it must appear as the first entry in the file. The JAR File Specification[2] describes several other optional attributes which can appear, but applications and add-on frameworks (such as OSGi) are free to define additional headers.

All of the other attributes shown here are defined by OSGi, but most are optional. To create a valid bundle, only the `Bundle-SymbolicName` attribute is mandatory.

2.2 OSGi Development Tools

In theory, one does not need any tools for building OSGi bundles beyond the standard Java tools: `javac` for Java source code compilation, `jar` for packaging, and a straightforward text editor for creating the `MANIFEST.MF`.

However very few Java programmers work with such basic tools because they require lots of effort to use, both in repeatedly typing long command lines and in remembering to execute all the steps in the right sequence. In practice we use build tools like Ant or Maven, and IDEs like Eclipse, NetBeans or IntelliJ. The same is true when developing for OSGi.

Likewise, directly editing the `MANIFEST.MF` file and constructing bundles with the `jar` command is burdensome. In particular the format of the `MANIFEST.MF` file is designed primarily for efficient processing by the JVM rather than for manual editing, and therefore it has some unusual rules that make it hard to work with directly. For example, each line is limited to 72 *bytes* (not characters!) of UTF-8 encoded data¹. Also OSGi’s syntax requires some strings to be repeated many times within the same file. Again, this is done for speed of automated processing rather than for convenient editing by humans².

There are tools that can make the task of developing and building bundles easier and less error-prone, so in this section we will review some of the tools available.

2.2.1 Eclipse Plug-in Development Environment

Eclipse is built on OSGi. For the Eclipse platform and community to grow, it was (and still is) essential to make it easy to produce new plug-ins that extend the functionality of Eclipse. Therefore, Eclipse provides a rich set of tools for building plug-ins. Taken together these tools are called the Plug-in Development Environment, usually abbreviated to PDE, which works as a layer on top of the Java Development Tools (JDT). Both PDE and JDT are included with the “Classic” Eclipse SDK download package, but PDE is *not* included with the “Eclipse IDE for Java Developers” download.

Since Eclipse plug-ins are just OSGi bundles³, the PDE can be used for OSGi

¹A single UTF-8 character is represented by between one and six bytes of data.

²It is the author’s opinion that these goals — efficient processing by machines, and ease of editing by humans — are fundamentally at odds. Formats convenient for humans to write are hard for machines to parse, and *vice versa*. XML attempts to be convenient for both, but manages to be difficult for humans to read and write while *still* being inefficient for machines to parse and generate.

³The term “plug-in” has always been used by Eclipse since its first version, which was not based on OSGi. Today there is no difference at all between a “plug-in” and a bundle. More details about the relationship between Eclipse and OSGi can be found in Chapter ??

development. In fact it is one of the simplest ways to very quickly and easily create new bundles.

However, we will not use PDE for the examples in this book, for a number of reasons:

- Most importantly, using PDE would force you to use Eclipse to work through the examples. Eclipse is a great IDE and the author's first choice for Java development, but OSGi should be accessible to users of other IDEs — such as NetBeans and IntelliJ — and even to those who prefer to use Emacs or Vim⁴!
- Second, PDE is a highly interactive and graphical tool. Unfortunately this is inconvenient for pedagogical purposes, since it is difficult to describe the steps needed to complete a complex operation without filling up the book with numerous large screenshots.
- Finally, in the author's opinion PDE encourages some “anti-patterns”, or practices which go against the recommended OSGi approach⁵.

2.2.2 Bnd

Bnd^[10] is a command-line tool, developed by Peter Kriens, for building OSGi bundles. Compared to PDE it may seem primitive, because it offers no GUI, instead using plain text properties files and instructions issued on the command line. However the core of the tool is very elegant and powerful, and it combines well with standard and familiar build tools such as ANT.

Listing 2.2 A Typical Bnd Descriptor File

```
# sample.bnd
Private-Package: org.example
Bundle-Activator: org.example.MyActivator
```

Listing 2.2 shows an example of a typical Bnd descriptor file. This looks very similar to a bundle manifest file, so to avoid confusion we will prefix bnd files with a comment indicating the file name, e.g. `# sample.bnd`. These files are much easier to edit than a `MANIFEST.MF` thanks to the following differences:

- As an alternative to the colon+space separator for the name/value pairs, an equals sign (=) can be used with or without surrounding spaces.
- There is no line length limit.

⁴Though I draw the line at Notepad.

⁵Specifically I feel that PDE has very weak support for the `Import-Package` form of dependency declaration, thus encouraging the use of `Require-Bundle`. The meaning of these terms is explained in Chapter 3.

- Line continuations, i.e. long strings broken over several lines to make them more readable, are indicated with a backslash character (\) at the end of the continued line.
- Shortcuts and wildcards are available in the `bnd` syntax for entering data that, in the underlying `MANIFEST.MF` syntax, would need to be repeated.

This descriptor file is used by `bnd` to generate not just the `MANIFEST.MF` but the JAR itself. The descriptor tells `bnd` both how to generate the `MANIFEST.MF` and also what the contents of the JAR should be. In this case the `Private-Package` instruction tells `bnd` that the JAR should contain the classes and resources found in the `org.example` package. It will use its classpath — which can be supplied either on the command line, or through the properties of the Ant task — to source the contents of the specified package.

2.3 Installing a Framework

In the introduction were listed four open source implementations of the OSGi standard. As was mentioned, all the sample code in this book is intended to run on all of the OSGi Release 4 implementations — i.e. Equinox, Knopflerfish and Felix, but not Concierge which is a Release 3 implementation⁶ — however each framework has its own idiosyncratic way to be launched and controlled. Sadly it would be impractical to show instructions for working with all three frameworks, so we must choose one, but fortunately the differences are slight enough that almost all the instructions can still be followed on another framework simply using the cross-reference of commands in Appendix ??.

We will mainly work with Apache Felix. It has a few advantages: in particular it is a quite compact R4 implementation, coming in at around 325K for the core framework JAR versus Equinox's 960K (Knopflerfish is also around 325K). Furthermore, Felix is slightly stricter in its interpretation of the OSGi specification, so a bundle that works on Felix is more or less guaranteed to work on Equinox and Knopflerfish. Finally, it is easy to configure and to embed into other Java applications. None of these reasons are particularly strong and they are not criticisms of the other frameworks; essentially the decision is little more than arbitrary.

The main download page for Felix, at <http://felix.apache.org/> unfortunately offers Felix in a fragmented series of downloads. The purpose of this is to reduce the minimal download size, but for convenience we would like to get not just the core binary framework but also the source code for Felix and the OSGi APIs. Source code is particularly useful when using an IDE, as it

⁶In fact many code samples *will* work on Concierge/OSGi R3, however in general R4 is assumed.

allows the IDE to infer more information about the API when it offers content assistance. Therefore a slightly modified download containing this addition is available from the author's website:

- <http://neilbartlett.name/downloads/felix-1.0.3-withsrc.zip>

Or if you prefer, you can download the standard binary package (with no source code) from the original Felix download page:

- <http://felix.apache.org/site/downloads.cgi>

Once you have downloaded either one of these packages, it can be extracted to a location of your choosing. After decompressing you will have a directory named `felix-1.0.3`, which we will refer to from now on as `FELIX_HOME`.

2.4 Setting up Eclipse

Section 2.2.1 discussed several problems with the Eclipse Plug-in Development Environment. However, in the examples we *will* still use the Eclipse IDE, because Java development is far easier and more productive in an IDE than in a plain text editor, and because Eclipse is by far the most popular IDE for Java. But rather than using PDE, a tool which is available only in Eclipse, we will use the basic Java Development Tooling (JDT), which has direct parallels in other Java IDEs. Therefore, although the instructions given are specific to Eclipse, they can be directly translated to those other IDEs. Please check the documentation for your preferred IDE: most include a section on how to translate from Eclipse concepts.

Before creating any projects, we will first define a “User Library” for the Felix framework, which will help us to reference Felix from many individual projects. To do this open the system preferences dialogue by selecting **Window** → **Preferences** (Mac users: **Eclipse** → **Preferences**) and navigating the tree on the left edge to **Java** → **Build Path** → **User Libraries**. Then click the **New** button and type the library name “Felix”. Next, with the Felix library selected, click **Add JARs** and add `bin/felix.jar` from the Felix directory (`FELIX_HOME`).

If you chose to download the source package, you need to inform Eclipse of the location of the source code. Expand the `felix.jar` entry in the User Library list, select “Source attachment” and click **Edit**. Click **External File** and browse to the file `src.zip` under `FELIX_HOME`.

You should now have something that looks like Figure 2.1.

Now we can create a new Java Project by selecting **File** → **New** → **Java Project**. Eclipse will show a wizard dialogue box. Enter the project name

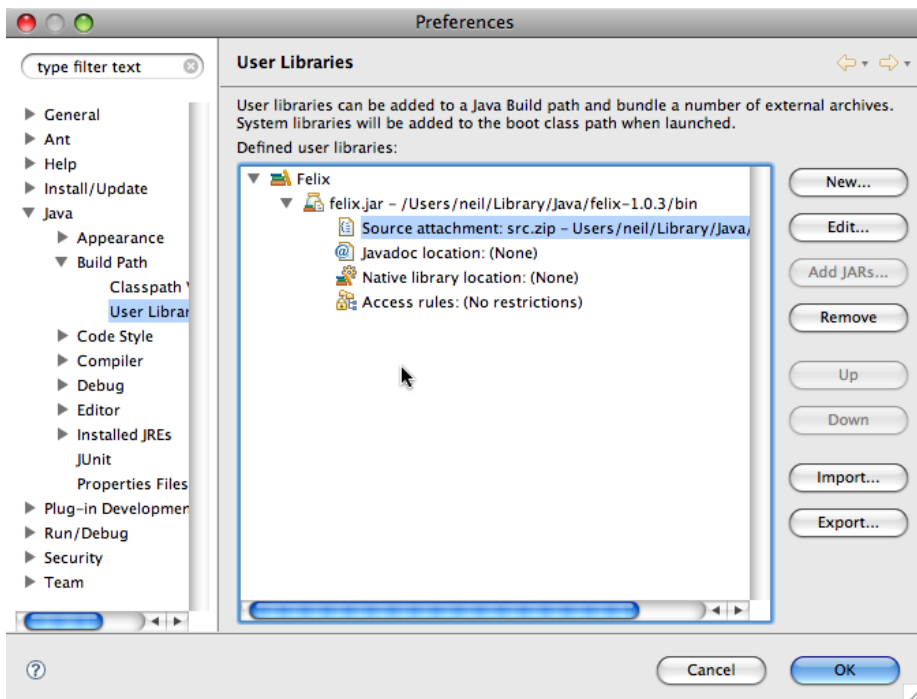


Figure 2.1: Adding Felix as a User Library in Eclipse

“OSGi Tutorial” and accept all the other defaults as given. Click **Next** to go to the second page of the wizard. Here we can add the Felix library by selecting the **Libraries** tab and clicking **Add Library**. A sub-wizard dialogue pops up: select **User Library** and click **Next**, then tick the box next to Felix and click **Finish**. Back in the main wizard, which should now look like Figure 2.2, we can click **Finish**, and Eclipse will create the project.

The next step is to copy some configuration for Felix and a few bundles we will be using. This step might seem a little bit strange at first, but it will become clear later why we are doing it this way. For now, copy two directories from the Felix distribution directory: **conf** and **bundles**. Your project should now look like Figure 2.3.

Before running Felix, let’s change the default logging level to something a little less noisy. Open the file **conf/config.properties** and search for the property **felix.log.level**. Change the value of the property from 4 to 1.

2.5 Running Felix

We’re now going to run Felix using the Eclipse launcher. We could run Felix from a shell or command prompt window, but if we run under Eclipse we can easily switch to debugging mode when things go wrong.

From the main menu, select **Run** → **Run Configurations** (or **Open Run Dialog** in older Eclipse versions). On the left hand side select “Java Application” and click the “New” button. Change the Name field at the top to “Felix”. Ensure the **Project** field is set to the “OSGi Tutorial” project, and then click the **Search** button next to the **Main class** field. A short list of “main” classes should pop up: choose **org.apache.felix.main.Main**.

Now click **Run**. Felix will print the following:

```
Welcome to Felix.  
-----  
Enter profile name:
```

“Profiles” are Felix’s mechanism for keeping different instances of the framework distinct. In general, OSGi frameworks are required to persist most of their state when shut down, and restore that state when restarted. This is a powerful feature, but it can sometimes be inconvenient if we have different active projects or development tasks that we need to switch between. Therefore Felix allows us to choose a profile when it is started, and it only restores the persisted state for the profile chosen. New profiles are created on the fly if one does not exist yet with the name provided.

Enter the profile name “tutorial” and hit enter. Felix will now start, and will print its standard shell prompt, which is `->` . We are now ready to run some

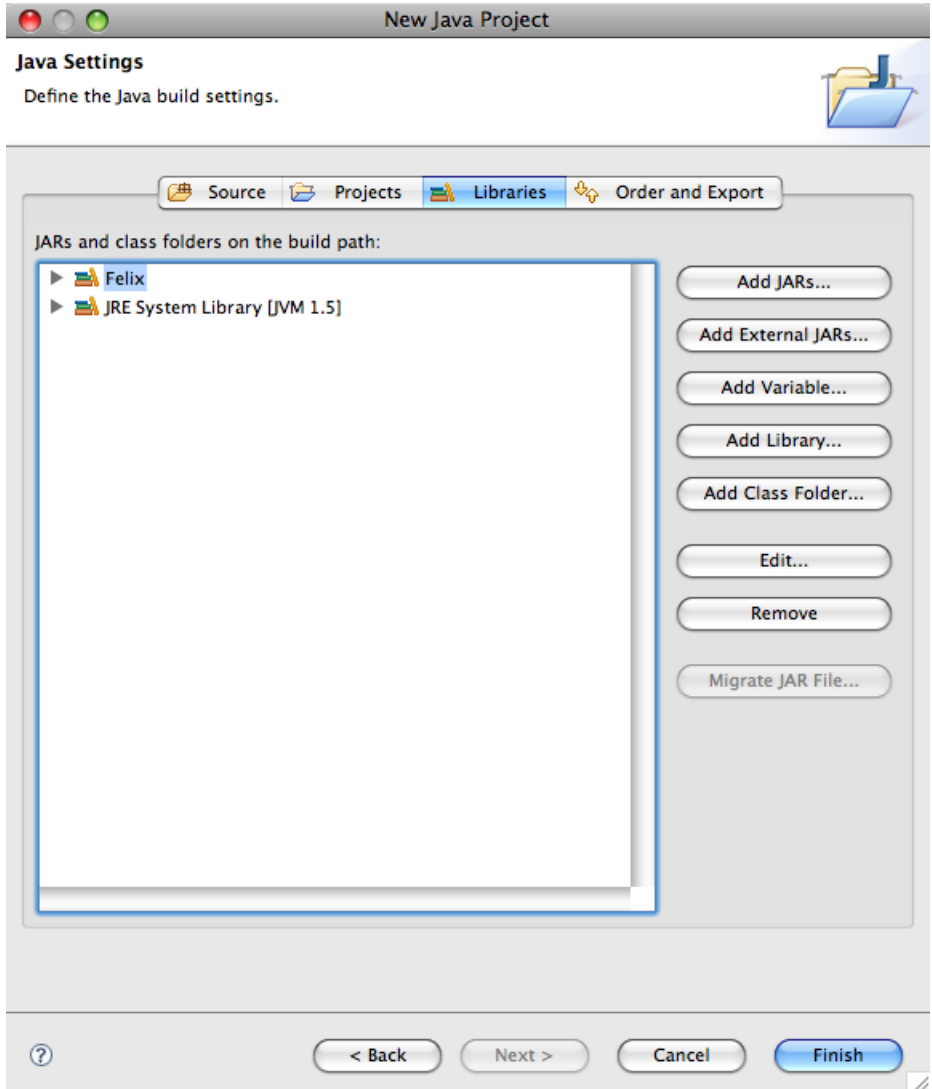


Figure 2.2: Creating a new Java project in Eclipse: adding the Felix library

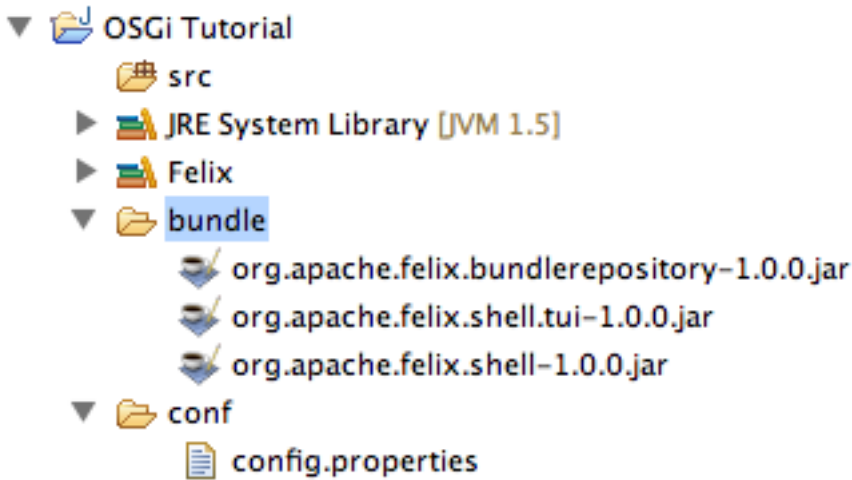


Figure 2.3: A New OSGi Project, Ready to Start Work

commands. The most frequently used command is `ps`, which prints the list of currently installed bundles along with their state. Let’s try running `ps` now:

```

-> ps
START LEVEL 1
  ID   State      Level  Name
[  0] [Active] [  0] System Bundle (1.0.1)
[  1] [Active] [  1] Apache Felix Shell Service (1.0.0)
[  2] [Active] [  1] Apache Felix Shell TUI (1.0.0)
[  3] [Active] [  1] Apache Felix Bundle Repository (1.0.0)

```

The so-called “System Bundle”, which always has the bundle ID of 0, is the framework itself. The other three bundles are the ones we copied over from the Felix distribution directory. The “Shell Service” and “Shell TUI” (Textual User Interface) bundles are providing the interactive shell we are using at the moment. The “Bundle Repository” bundle provides an interface to download additional bundles from the OSGi Bundle Repository — we will look at that feature in Chapter ??.

For the moment, there is not much more we can do with Felix until we install some more interesting bundles. However you may at this point wish to explore the commands available by typing `help`. When bored with this, don’t shutdown Felix just yet; leave it running for now.

2.6 Installing bnd

Peter Kriens’ `bnd` is an ingenious little tool. It is packaged as a single JAR file, yet it is simultaneously:

- a standalone Java program, which can be invoked with the `java -jar` command;
- an Eclipse plug-in;
- an Ant task;
- a Maven plug-in.

We will install it as an Eclipse plug-in, but we will also be using it standalone and as an Ant task soon. First download it from:

<http://www.aqute.biz/Code/Download#bnd>

then take a copy and save it into the `plugins` directory underneath your Eclipse installation directory⁷. Put another copy into the project directory in Eclipse. Finally, restart Eclipse with the `-clean` command line parameter.

Next we will configure Eclipse to use its standard Java “properties” file editor to open files ending in the `.bnd` extension. To do this, open the system preferences and navigate to **General** → **Editors** → **File Associations**. Click the **Add** button next to the upper list and type `*.bnd`. Now click the **Add** button next to the lower list and select **Internal Editors** → **Properties File Editor**. Click **OK** to close the preferences window.

2.7 Hello, World!

In keeping with long-standing tradition, our first program in OSGi will be one that simply prints “Hello, World” to the console. However, most such programs immediately exit as soon as they have printed the message. We will embrace and extend the tradition with our first piece of OSGi code: since OSGi bundles have a concept of lifecycle, we will not only print “Hello” upon start-up but also “Goodbye” upon shutdown.

To do this we need to write a bundle activator. This is a class that implements the `BundleActivator` interface, one of the most important interfaces in OSGi.

Bundle activators are very simple. They have two methods, `start` and `stop`, which are called by the framework when the bundle is started and stopped respectively. Our “Hello/Goodbye, World!” bundle activator is as simple as the class `HelloWorldActivator` as shown in Listing 2.3.

There’s not a lot to explain about this class, so let’s just get on and build it. We need a bnd descriptor file, so create a file at the top level of the project called `helloworld.bnd`, and copy in the following contents:

⁷Or the `dropins` directory if you are using Eclipse 3.4 or later.

Listing 2.3 Hello World Activator

```

1 package org.osgi.tutorial;
2
3 import org.osgi.framework.BundleActivator;
4 import org.osgi.framework.BundleContext;
5
6 public class HelloWorldActivator implements BundleActivator {
7     public void start(BundleContext context) throws Exception {
8         System.out.println("Hello, World!");
9     }
10
11    public void stop(BundleContext context) throws Exception {
12        System.out.println("Goodbye, World!");
13    }
14 }

```

Listing 2.4 Bnd Descriptor for the Hello World Activator

```

# helloworld.bnd
Private-Package: org.osgi.tutorial
Bundle-Activator: org.osgi.tutorial.HelloWorldActivator

```

This says that the bundle we want to build consists of the specified package, which is private (i.e. non-exported), and that the bundle has an activator, namely the class we just wrote.

If you installed `bnd` as an Eclipse plug-in, you can now right-click on the `helloworld.bnd` file and select **Make Bundle**. As a result, `bnd` will generate a bundle JAR called `helloworld.jar`. However, if you are *not* using Eclipse or the `bnd` plug-in for Eclipse, you will have to build it manually: first use the `javac` compiler to generate class files, then use the command

```
java -jar bnd.jar build -classpath classes helloworld.bnd
```

where `classes` is the directory containing the class files. Alternatively you could try the ANT build script given in Appendix B.

Now we can try installing this bundle into Felix, which you should still have running from the previous section (if not, simply start it again). At the Felix shell prompt, type the command:

```
install file:helloworld.jar
```

The `install` command always takes a URL as its parameter rather than a file name, hence the `file:` prefix which indicates we are using the local file URI scheme[11]. Strictly, this URL is invalid because it is relative rather than absolute, however Felix allows you, in the interests of convenience, to provide a path relative to the current working directory. Therefore since the working directory is the project directory, we need only the file name in this case.

Felix will respond with the bundle ID that it has assigned to the bundle:

```
Bundle ID: 4
```

That ID will be used in subsequent commands to manipulate the bundle. It's worth noting at this point that the bundle ID you get *may* be different to what you see in this text. That's more likely to happen later on, but it's important to be aware of your actual bundle IDs and try not to copy slavishly what appears in these examples.

Let's take a quick look at the bundle list by typing `ps`. It should look like this:

```
START LEVEL 1
  ID      State      Level  Name
[  0] [Active  ] [  0] System Bundle (1.0.3)
[  1] [Active  ] [  1] Apache Felix Shell Service (1.0.0)
[  2] [Active  ] [  1] Apache Felix Shell TUI (1.0.0)
[  3] [Active  ] [  1] Apache Felix Bundle Repository (1.0.2)
[  4] [Installed] [  1] helloworld (0)
```

Now the moment of truth: start the new bundle by typing `start 4`. We should see:

```
-> start 4
Hello , World!
```

What does the bundle list look like now if we type `ps`? We'll ignore the rest of the listing, as it hasn't changed, and focus on our bundle:

```
[  4] [Active  ] [  1] helloworld (0)
```

The bundle is now in the "Active" state, which makes sense since we have explicitly started it. Now stop the bundle with `stop 4`:

```
-> stop 4
Goodbye , World!
```

This works as expected. What does the bundle list look like now?

```
[  4] [Resolved ] [  1] helloworld (0)
```

This is interesting. When we started the bundle it changed its state from Installed to Active, but when we stopped the bundle it didn't go back to Installed: instead it went to a new state "Resolved". What's going on here? Time to look at the lifecycle of a bundle.

2.8 Bundle Lifecycle

It was mentioned that OSGi bundles have a lifecycle, but what exactly is that lifecycle? How and why do bundles move from one state to another? Figure 2.4 shows the full lifecycle.

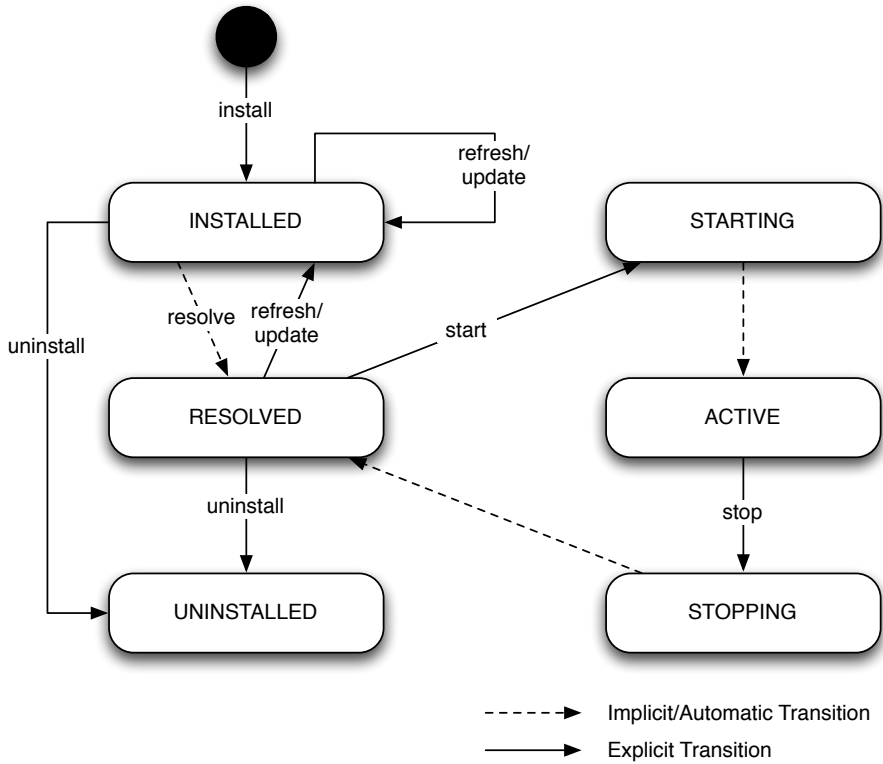


Figure 2.4: Bundle Lifecycle

It's worth spending some time to trace the path through this diagram that was taken by our "Hello/Goodbye, World" bundle. Like all bundles, it started at the black circle, the entry point of the diagram, before we installed it by executing the `install` command. At that point it entered the `INSTALLED` state.

Next we executed the `start` command and it appeared to transition directly to the `ACTIVE` state, although the diagram shows no direct link between those two states. Strictly speaking, bundles can only be started when they are in `RESOLVED` state, however when we attempt to start an `INSTALLED` bundle, the framework simply attempts to resolve it first before proceeding to start it.

`RESOLVED` means that the bundle's constraints have all been satisfied. In other words:

- The Java execution environment — e.g. CDC Foundation, Java 5, etc — matches or exceeds what was specified by the bundle;
- The imported packages of the bundle are available and exported with the right version range by other `RESOLVED` bundles, or bundles that can be `RESOLVED` at the same time as this bundle;
- The required bundles of the bundle are available and `RESOLVED`, or can be `RESOLVED`.

Once those constraints are satisfied, a bundle can be resolved, i.e. moved from the `INSTALLED` state to the `RESOLVED` state. If any of those constraints becomes unsatisfied — e.g. another bundle providing an imported package has been removed — then the bundle must be unresolved, i.e. moved from `RESOLVED` back to `INSTALLED`.

When we executed the `start` command on an `INSTALLED` bundle, the framework noticed that it first needed to attempt to resolve the bundle. It did so immediately, because in this case our bundle was so simple that it didn't have any constraints that needed to be satisfied. So the transition to `RESOLVED` state was automatic, and quickly followed by a transition to `STARTING` state.

`STARTING` simply means that the bundle is in the process of being activated. For example, the framework is currently calling the `start` method of its bundle activator. Usually the `STARTING` state is over in the blink of an eye, and you would not be able to see any bundles in this state on the bundle list.

Once the `STARTING` process is complete, the bundle reaches `ACTIVE` state. In this state the bundle is not necessarily actively running any code. The activity that happens during the `ACTIVE` state of a bundle generally depends on whatever was kicked off by the `start` method of the bundle activator.

When the `stop` command is executed, the bundle transitions to `STOPPING` state while the framework calls the `stop` method of its bundle activator. This

is a chance for the bundle to terminate and clean up anything that it created during the STARTING phase. STOPPING is also a transient state which exists for a very short period of time before the bundle returns to RESOLVED state.

Finally we can choose to un-install the bundle, although we did not do it in the example above. The diagram is slightly misleading: although the UNINSTALLED state is shown here, we can never see a bundle in that state, and an UNINSTALLED bundle cannot transition to any other state. Even if we reinstall the same bundle JAR file, it will be considered a different bundle by the framework, and assigned a new bundle ID.

2.9 Incremental Development

Development tends to work best as an incremental process: we prefer to make small changes and test them quickly, because we can find and correct our mistakes and misunderstandings sooner. If we write large blocks of code before testing them, there's a good chance we will have to rewrite them in full, so incremental development wastes less effort.

In some programming languages such as Scheme, development is centred around a “Read-Evaluate-Print Loop” (REPL) or interactive shell which gives us quick feedback about small changes to our code. This style of development has always been difficult in Java, as in other compiled languages, since the code must be built and deployed before it can be run, and redeployment usually implied restarting the JVM. In extreme cases such as J2EE development we might have to run a five-minute build script and then spend another five minutes restarting our application server.

OSGi can help. Although we are unlikely to ever be able to make the development cycle for Java as tight as that for Scheme, the modularity of our code and the ability to install and un-install individual bundles on the fly means that we don't need to “rebuild the world” when we make a small change, and we don't need to restart anything except the individual bundle.

Suppose we make a change to the “Hello, World!” bundle, for example we would like to print the messages in French instead of English. We can change the code as necessary and rebuild just this bundle by right-clicking on `hello-world.bnd` and selecting `Make Bundle`.

Back in the Felix shell, we could choose to un-install the old `helloworld` bundle and install it again, but that would result in a new bundle with a new bundle ID. In fact we can simply update the existing bundle:

```
-> update 4
-> start 4
Bonjour le Monde !
```

```
-> stop 4
Au revoir !
```

Note that we didn't have to tell Felix where to update the bundle from: it remembers the URL from which it was originally installed, and simply re-reads from that URL. This still works even if Felix has been shutdown and restarted since the bundle was installed. Of course sometimes we wish to update from a new location, in which case we can pass that location as a second parameter to the `update` command.

2.10 Interacting with the Framework

Taking a look back at the code for `HelloWorldActivator`, we see that something else happens in the bundle activator. When the framework calls our activator's `start` and `stop` methods, it passes in an object of type `BundleContext`.

The bundle context is the “magic ticket” we need to interact with the OSGi framework from our bundle code. *All* interaction with the framework goes through the context, and the *only* way to access a bundle context is to implement an activator and have the framework give it to us⁸.

So, what sort of things can we do with the framework via `BundleContext`? Here is an (incomplete) list:

- Look up system-wide configuration properties;
- Find another installed bundle by its ID;
- Obtain a list of all installed bundles;
- Introspect and manipulate other bundles programmatically: start them, stop them, un-install them, update them, etc;
- Install new bundles programmatically;
- Store or retrieve a file in a persistent storage area managed by the framework;
- Register and unregister bundle listeners, which tell us when the state of any bundle in the framework changes;
- Register and unregister service listeners, which tell us when the state of any service in the framework changes (services and service listeners are the subject of Chapter 4);

⁸In fact this is a small lie. There is another way to get a bundle context, which is discussed in Section 8.6 of Chapter 8, but this is an advanced technique and would only confuse matters at this stage.

- Register and unregister framework listeners, which tell us about general framework events.

Time for another example. Suppose we are *very* interested in the total number of bundles currently installed. We would like a bundle that lets us know if a bundle is installed or un-installed, along with the new total number of bundles. It should also print the total number when it starts up.

One way to approach this would be to attempt to track the running total of bundles ourselves, by first retrieving the current total when our bundle is started, and incrementing and decrementing the total as we are notified of bundles being installed and un-installed. However, that approach quickly runs into some tricky multi-threading issues. It is also unnecessary since the framework tracks all bundles anyway, so we can simply ask the framework for the current total each time we know that the total has changed. This approach is shown in Listing 2.5.

This bundle activator class is also a bundle listener: it implements both interfaces. When it starts up, it registers itself as a bundle listener and then prints the total number of bundles. Incidentally, we *must* do it that way around: if we first printed the total and then registered as a listener, there would be a potential problem, as follows. Suppose the bundle count were 10, and then a new bundle happened to be installed just before our listener started working. That is, a bundle is installed between our calls to `getBundles` and `addBundleListener`. We would get no message regarding the new bundle because we missed the event, so we would still think the bundle count is 10 when in fact it is 11. If another bundle were to be installed later, we would see the following confusing series of messages:

```
There are currently 10 bundles
Bundle installed
There are currently 12 bundles
```

By registering the listener first, we can be sure not to miss any events, but there is a price: we may get duplicate messages if a bundle is installed or un-installed between registering the listener and printing the total. In other words we may see this series of messages:

```
Bundle installed
There are currently 11 bundles
There are currently 11 bundles
```

The repetition may be annoying, but it is substantially less confusing than the apparent “jump” from ten bundles to twelve. Listing 2.6 shows the `bnd` descriptor that we need to build the bundle, `bundlecounter.bnd`.

To test the bundle counter, install and start it as before and then test it by un-installing and re-installing the `helloworld` bundle.

Listing 2.5 Bundle Counter Activator

```

1 package org.osgi.tutorial;

3 import org.osgi.framework.BundleActivator;
4 import org.osgi.framework.BundleContext;
5 import org.osgi.framework.BundleEvent;
6 import org.osgi.framework.BundleListener;

8 public class BundleCounterActivator implements BundleActivator,
9     BundleListener {

11     private BundleContext context;

13     public void start(BundleContext context) throws Exception {
14         this.context = context;

16         context.addBundleListener(this); //1
17         printBundleCount();           //2
18     }

20     public void stop(BundleContext context) throws Exception {
21         context.removeBundleListener(this);
22     }

24     public void bundleChanged(BundleEvent event) {
25         switch (event.getType()) {
26             case BundleEvent.INSTALLED:
27                 System.out.println("Bundle installed");
28                 printBundleCount();
29                 break;
30             case BundleEvent.UNINSTALLED:
31                 System.out.println("Bundle uninstalled");
32                 printBundleCount();
33                 break;
34         }
35     }

37     private void printBundleCount() {
38         int count = context.getBundles().length;
39         System.out.println("There are currently " + count + " bundles");
40     }
41 }

```

Listing 2.6 Bnd Descriptor for the Bundle Counter

```

# bundlecounter.bnd
Private-Package: org.osgi.tutorial
Bundle-Activator: org.osgi.tutorial.BundleCounterActivator

```

2.11 Starting and Stopping Threads

One of the things that we can do from the `start` method of a bundle activator is start a thread. Java makes this very easy, but the part that we must worry about is cleanly stopping our threads when the bundle is stopped.

Listing 2.7 Heartbeat Activator

```
1 package org.osgi.tutorial;
2
3 import org.osgi.framework.BundleActivator;
4 import org.osgi.framework.BundleContext;
5
6 public class HeartbeatActivator implements BundleActivator {
7
8     private Thread thread;
9
10    public void start(BundleContext context) throws Exception {
11        thread = new Thread(new Heartbeat());
12        thread.start();
13    }
14
15    public void stop(BundleContext context) throws Exception {
16        thread.interrupt();
17    }
18 }
19
20 class Heartbeat implements Runnable {
21
22    public void run() {
23        try {
24            while (!Thread.interrupted()) {
25                Thread.sleep(5000);
26                System.out.println("I'm still here.");
27            }
28        } catch (InterruptedException e) {
29            System.out.println("I'm going now.");
30        }
31    }
32 }
```

The code in Listing 2.7 shows a simple “heartbeat” thread that prints a message every five seconds. In this case, we can use Java’s interruption mechanism, which is a simple boolean status that can be set on a thread by calling the `interrupt` method. Unfortunately it is not always so easy to wake up a thread and ask it to stop, as we will see in Chapter 6.

2.12 Manipulating Bundles

The OSGi framework gives us quite a lot of programmatic control over other bundles. Let’s look at an example in which we tie the lifecycle of a “target” bundle to its source file in the filesystem.

The proposed scenario is as follows: we have a bundle — our original “Hello, World!” bundle, say — that is changing frequently (we just can’t decide which language those messages should be in!). Every time the bundle is rebuilt we could simply type the `update` command, but even this can be a hassle after a while. It can be easy to forget to update a bundle, and then wonder why the code you just wrote isn’t working. Therefore we would like some kind of “manager” bundle that will continually monitor the `helloworld.jar` file and execute the update for us when the file changes.

The code in Listing 2.8 does exactly that. It uses a polling loop, like the heartbeat example from the last section, but on each beat it checks whether the file `helloworld.jar` is newer, according to its last-modified timestamp, than the corresponding bundle. If it is, then it updates the bundle, causing it to re-load from the file. If the bundle is up to date, or if either the bundle or the file do not exist, then it does nothing.

2.13 Exercises

1. Write a bundle that periodically checks the contents of a directory in the filesystem. Whenever a new file with the extension `.jar` appears in that directory, attempt to install it as a bundle, but only if a corresponding bundle is not already present.
2. Extend the last exercise by checking for deleted files. If a file corresponding to an installed bundle is deleted, then un-install that bundle.
3. Finally, extend your bundle to detect changes in all files in the directory that correspond to bundles, and update those bundles as necessary.

Listing 2.8 Hello Updater Activator

```
1 package org.osgi.tutorial;
2
3 import java.io.File;
4
5 import org.osgi.framework.Bundle;
6 import org.osgi.framework.BundleActivator;
7 import org.osgi.framework.BundleContext;
8 import org.osgi.framework.BundleException;
9
10 public class HelloUpdaterActivator implements BundleActivator {
11
12     private static final long INTERVAL = 5000;
13     private static final String BUNDLE = "helloworld.jar";
14
15     private final Thread thread = new Thread(new BundleUpdater());
16     private volatile BundleContext context;
17
18     public void start(BundleContext context) throws Exception {
19         this.context = context;
20         thread.start();
21     }
22
23     public void stop(BundleContext context) throws Exception {
24         thread.interrupt();
25     }
26
27     protected Bundle findBundleByLocation(String location) {
28         Bundle[] bundles = context.getBundles();
29         for (int i = 0; i < bundles.length; i++) {
30             if (bundles[i].getLocation().equals(location)) {
31                 return bundles[i];
32             }
33         }
34
35         return null;
36     }
37
38     private class BundleUpdater implements Runnable {
39         public void run() {
40             try {
41                 File file = new File(BUNDLE);
42                 String location = "file:" + BUNDLE;
43                 while (!Thread.currentThread().isInterrupted()) {
44                     Thread.sleep(INTERVAL);
45                     Bundle bundle = findBundleByLocation(location);
46                     if (bundle != null && file.exists()) {
47                         long bundleModified = bundle.getLastModified();
48                         long fileModified = file.lastModified();
49                         if (fileModified > bundleModified) {
50                             System.out.println("File is newer, updating");
51                             bundle.update();
52                         }
53                     }
54                 }
55             } catch (InterruptedException e) {
56                 System.out.println("I'm going now.");
57             } catch (BundleException e) {
58                 System.err.println("Error updating bundle");
59                 e.printStackTrace();
60             }
61         }
62     }
63 }
```


3 Bundle Dependencies

In the last chapter we created some simple bundles and showed how those bundles can interact with the framework. In this chapter, we will start to look at how bundles can interact with each other. In particular we will see how to manage dependencies between bundles.

As discussed in the Introduction, managing dependencies is the key to achieving modularity. This can be a big problem in Java — and in many other languages as well, since few provide the kind of module systems that are needed to build large applications. The default module system in Java (and it is a stretch to call it that) is the JAR-centric “classpath” model, which fails mainly because it does not manage dependencies, but instead leaves them up to chance.

What is a dependency? Simply, it is a set of assumptions made by a program or block of code about the environment in which it will run.

For example, a Java class may assume that it is running on a specific version of the Java VM, and that it has access to a specific library. It therefore depends on that Java version and that library. However those assumptions are *implicit*, meaning we don’t know that they exist until the code first runs in an environment in which they are false, and an error occurs. Suppose a Java class depends on the Apache Log4J library — i.e. it assumes the Log4J JAR is available on the classpath — and then it is run in an environment where that is false (Log4J is *not* on the classpath). It will produce errors at runtime such as `ClassNotFoundException` or `NoClassDefFoundError`. Therefore we have no real idea whether the class will work in any particular environment except by trying it out, and even if it initially appears to work, it may fail at a later time when a particular execution path is followed for the first time and that path exposes a new, previously unknown, dependency.

OSGi takes away the element of chance, by managing dependencies so that they are *explicit*, *declarative* and *versioned*.

Explicit A bundle’s dependencies are in the open for anybody to see, rather than hidden in a code path inside a class file, waiting to be found at runtime.

Declarative Dependencies are specified in a simple, static, textual form for easy inspection. A tool can calculate which set of bundles are required to

satisfy the dependencies of a particular bundle without actually installing or running any of them.

Versioned Libraries change over time, and it is not enough to merely depend on a library without regard to its version. OSGi therefore allows all inter-bundle dependencies to specify a version range, and even allows for multiple versions of the same bundle to be present and in use at the same time.

3.1 Introducing the Example Application

Rather than working on multiple small example pieces of code which never amount to anything interesting, for the remainder of this book we will start to build up a significant example application. The example will be based around the concept of a “message centre”, for processing and displaying messages from multiple sources.

Thanks to the internet, we are today bombarded by messages of many kinds. Email is an obvious example, but also there are the blogs that we subscribe to through RSS or ATOM feeds; SMS text messages; “microblogging” sites such as Twitter[12] or Jaiku[13]; IM systems such as AOL, MSN or IRC; and perhaps for professionals in certain fields, Reuters or Bloomberg newswire feeds and market updates. Sadly we still need to flip between several different applications to view and respond to these messages. Also there is no coherent way to apply rules and automated processing to all of our inbound messages. For example, I would like a way to apply my spam filters, which do a reasonable job for my email, to my SMS text messages, as spam is increasingly a problem in that medium.

So, let’s build a message centre application with two principal components:

- A graphical “reader” tool for interactively reading messages.
- An agent platform for running automated processing (e.g. filtering) over inbound message streams. This component will be designed to run either in-process with the reader tool, or as a separate server executable.

3.2 Defining an API

To support multiple kinds of message sources, we need to build an abstraction over messages and mailboxes, so a good place to start is to think about what those abstractions should look like. In Java we would represent them as interfaces. Listing 3.1 contains a reasonable attempt to define a message in the most general way.

Listing 3.1 The Message Interface

```
1 package org.osgi.book.reader.api;
2
3 import java.io.InputStream;
4
5 public interface Message {
6
7     /**
8      * @return The unique (within this message's mailbox) message ID.
9      */
10    long getId();
11
12    /**
13     * @return A human-readable text summary of the message. In some
14     *         messaging systems this would map to the "subject" field.
15     */
16    String getSummary();
17
18    /**
19     * @return The Internet MIME type of the message content.
20     */
21    String getMIMEType();
22
23    /**
24     * Access the content of the message.
25     *
26     * @throws MessageReaderException
27     */
28    InputStream getContent() throws MessageReaderException;
29
30 }
```

Objects implementing this interface are really just message headers. The body of the message could be of any type: text, image, video, etc. We need the header object to tell us what type the body data is, and how to access it.

Listing 3.2 The Mailbox Interface

```

1 package org.osgi.book.reader.api;
2
3 public interface Mailbox {
4
5     public static final String NAME_PROPERTY = "mailboxName";
6
7     /**
8      * Retrieve all messages available in the mailbox.
9      *
10     * @return An array of message IDs.
11     * @throws MailboxException
12     */
13     long [] getAllMessages() throws MailboxException;
14
15     /**
16     * Retrieve all messages received after the specified message.
17     *
18     * @param id The message ID.
19     * @return An array of message IDs.
20     * @throws MailboxException
21     */
22     long [] getMessagesSince(long id) throws MailboxException;
23
24     /**
25     * Mark the specified messages as read/unread on the back-end
26     * message source, where supported, e.g. IMAP supports this
27     * feature.
28     *
29     * @param read Whether the specified messages have been read.
30     * @param ids An array of message IDs.
31     * @throws MailboxException
32     */
33     void markRead(boolean read, long [] ids) throws MailboxException;
34
35     /**
36     * Retrieve the specified messages.
37     *
38     * @param ids The IDs of the messages to be retrieved.
39     * @return An array of Messages.
40     * @throws MailboxException
41     */
42     Message [] getMessages(long [] ids) throws MailboxException;
43 }

```

Next we need a way to retrieve messages. The interface for a mailbox could look like Listing 3.2. We need a unique identifier to refer to each message, so we assume that an ID of type `long` can be generated or assigned by the mailbox implementation. We also assume that the mailbox maintains some temporal ordering of messages, and is capable of telling us about all the new messages available given the ID of the most recent message known about by the reader tool. In this way the tool can notify us only of new messages, rather than ones that we have already read.

Many back-end message sources, such as the IMAP protocol for retrieving email, support storing the read/unread state of a message on the server, allowing that state to be synchronized across multiple clients. So, our reader tool needs to notify the mailbox when a message has been read. In other protocols where the back-end message source does not support the read/unread status, this notification can be simply ignored.

Finally we need the code for the exceptions that might be thrown. These are shown in Listing 3.3.

Listing 3.3 Mailbox API Exceptions

```
1 package org.osgi.book.reader.api;
3 public class MessageReaderException extends Exception {
5     private static final long serialVersionUID = 1L;
7     public MessageReaderException(String message) {
8         super(message);
9     }
11    public MessageReaderException(Throwable cause) {
12        super(cause);
13    }
15    public MessageReaderException(String message, Throwable cause) {
16        super(message, cause);
17    }
19 }

1 package org.osgi.book.reader.api;
3 public class MailboxException extends Exception {
5     public MailboxException(String message) {
6         super(message);
7     }
9     public MailboxException(Throwable cause) {
10        super(cause);
11    }
13    public MailboxException(String message, Throwable cause) {
14        super(message, cause);
15    }
17 }
```

3.3 Exporting the API

Now let's package up these classes into an API bundle. Create a bnd descriptor named `mailbox_api.bnd`, as shown in Listing 3.4.

Listing 3.4 Bnd Descriptor for the Mailbox API

```
# mailbox_api.bnd
Export-Package: org.osgi.book.reader.api;version=1.0.0
```

Unlike the previous bnd descriptor files, this descriptor does not have either a `Private-Package` or a `Bundle-Activator` entry. The bundle we are building here does not need to interact with the OSGi framework, it simply provides API to other bundles, so we need the `Export-Package` directive, which instructs bnd to do *two* separate things:

- It ensures that the contents of the named packages are included in the output bundle JAR.
- It adds an `Export-Package` header to the `MANIFEST.MF` of the JAR.

From this descriptor, bnd will generate `mailbox_api.jar`, so let's take a look inside the JAR. On Windows you can use a tool such as WinZip[14] to peek inside; on UNIX platforms including Linux and Mac OS X there is a wider variety of tools, but the command-line tool `unzip` should be available on nearly all of them.

The file content listing of the JAR should be no surprise: we simply have the four API classes (two interfaces, two exceptions) in the normal Java package directory tree:

```
$ unzip -l mailbox_api.jar
Archive:  mailbox_api.jar
  Length      Name
-----
    302      META-INF/MANIFEST.MF
     0      org/
     0      org/osgi/
     0      org/osgi/book/
     0      org/osgi/book/reader/
     0      org/osgi/book/reader/api/
    379      org/osgi/book/reader/api/Mailbox.class
    677      org/osgi/book/reader/api/MailboxException.class
    331      org/osgi/book/reader/api/Message.class
    759      org/osgi/book/reader/api/MessageReaderException.class
-----
    2448      10 files
```

Let's take a look at the generated manifest. WinZip users can right click on `MANIFEST.MF` and select "View with Internal Viewer". UNIX users can run `unzip` with the `-p` ("pipe") switch:

```
$ unzip -p mailbox_api.jar META-INF/MANIFEST.MF
Manifest-Version: 1.0
Bundle-Name: mailbox_api
Created-By: 1.5.0_13 (Apple Inc.)
Import-Package: org.osgi.book.reader.api
Bundle-ManifestVersion: 2
Bundle-SymbolicName: mailbox_api
Tool: Bnd-0.0.223
```

```
Bnd-LastModified: 1198796713427
Export-Package: org.osgi.book.reader.api
Bundle-Version: 0
```

Here we see the `Export-Package` header, along with a few other headers which have been added by bnd. This header tells OSGi to make the named list of packages (in this case only one package) available to be imported by other bundles. We will take a look at how to import packages shortly.

As we saw in Section 1.2.5, plain Java makes all public classes in all packages of a JAR available to clients of that JAR, making it impossible to hide implementation details. OSGi uses the `Export-Package` header to fix that problem. In fact the default is reversed: in OSGi, all packages are hidden from clients unless they are explicitly named in the `Export-Package` header.

In bnd the general form of the `Export-Package` directive is a comma-separated list of patterns, where each pattern either a literal package name, or a “glob” (wildcard) pattern, or a negation introduced with the `!` character. For example the following directive instructs bnd to include every package it can find on the classpath except for those starting with “com.”:

```
Export-Package: !com.*, *
```

3.4 Importing the API

Now let’s write some code that depends on the API: an implementation of a mailbox and message types. For the sake of simplicity at this stage, this will be a mailbox that holds only a fixed number of hard-coded messages. The code will live in a new package, `org.osgi.book.reader.fixedmailbox`. First we write an implementation of the `Message` interface, giving us the class `StringMessage` in Listing 3.5, and next the `Mailbox` interface implementation, giving us `FixedMailbox`¹ in Listing 3.6.

Now, what should the bnd descriptor look like? In the last section we used the `Export-Package` directive to include the specified package in the bundle, and also export it. This time, we want to include our new package in the bundle but we don’t want to export it, so we use the `Private-Package` instruction:

```
# fixed_mailbox.bnd
Private-Package: org.osgi.book.reader.fixedmailbox
```

Unlike `Export-Package`, the `Private-Package` directive does not correspond to a recognized `MANIFEST.MF` attribute in OSGi. It is merely an instruction to bnd that causes the specified packages to be included in the bundle JAR —

¹You may be curious at this stage why methods of `FixedMailbox` are declared as `synchronized`. This is required to keep the implementation thread-safe when we extend it later, in Chapter 7

Listing 3.5 String Message

```
1 package org.osgi.book.reader.fixedmailbox;
2
3 import java.io.ByteArrayInputStream;
4 import java.io.InputStream;
5
6 import org.osgi.book.reader.api.Message;
7 import org.osgi.book.reader.api.MessageReaderException;
8
9 public class StringMessage implements Message {
10
11     private static final String MIME_TYPE_TEXT = "text/plain";
12
13     private final long id;
14     private final String subject;
15     private final String text;
16
17     public StringMessage(long id, String subject, String text) {
18         this.id = id;
19         this.subject = subject;
20         this.text = text;
21     }
22
23     public InputStream getContent() throws MessageReaderException {
24         return new ByteArrayInputStream(text.getBytes());
25     }
26
27     public long getId() {
28         return id;
29     }
30
31     public String getMIMEType() {
32         return MIME_TYPE_TEXT;
33     }
34
35     public String getSummary() {
36         return subject;
37     }
38 }
```

Listing 3.6 Fixed Mailbox

```
1 package org.osgi.book.reader.fixedmailbox;
2
3 import java.util.ArrayList;
4 import java.util.List;
5
6 import org.osgi.book.reader.api.Mailbox;
7 import org.osgi.book.reader.api.MailboxException;
8 import org.osgi.book.reader.api.Message;
9
10 public class FixedMailbox implements Mailbox {
11
12     protected final List<Message> messages;
13
14     public FixedMailbox() {
15         messages = new ArrayList<Message>(2);
16         messages.add(new StringMessage(0, "Hello", "Welcome to OSGi"));
17         messages.add(new StringMessage(1, "Getting Started",
18             "To learn about OSGi, read my book."));
19     }
20
21     public synchronized long[] getAllMessages() {
22         long[] ids = new long[messages.size()];
23         for (int i = 0; i < ids.length; i++) {
24             ids[i] = i;
25         }
26         return ids;
27     }
28
29     public synchronized Message[] getMessages(long[] ids)
30         throws MailboxException {
31         Message[] result = new Message[ids.length];
32         for (int i = 0; i < ids.length; i++) {
33             long id = ids[i];
34             if (id < 0 || id >= messages.size()) {
35                 throw new MailboxException("Invalid message ID: " + id);
36             }
37             result[i] = messages.get((int) id);
38         }
39         return result;
40     }
41
42     public synchronized long[] getMessagesSince(long id)
43         throws MailboxException {
44         int first = (int) (id + 1);
45         if (first < 0 || first >= messages.size()) {
46             throw new MailboxException("Invalid message ID: " + first);
47         }
48         long[] ids = new long[messages.size() - first];
49         for (int i = 0; i < ids.length; i++) {
50             ids[i] = i + first;
51         }
52         return ids;
53     }
54
55     public void markRead(boolean read, long[] ids) {
56         // Ignore
57     }
58 }
```

you can verify this by looking at the contents of the JAR. Listing 3.7 shows what the generated MANIFEST.MF should look like.

Listing 3.7 MANIFEST.MF generated from `fixed_mailbox.bnd`


```
Manifest-Version: 1.0
Bundle-Name: fixed_mailbox
Created-By: 1.5.0_13 (Apple Inc.)
Private-Package: org.osgi.book.reader.fixed
Import-Package: org.osgi.book.reader.api
Bundle-ManifestVersion: 2
Bundle-SymbolicName: fixed_mailbox
Tool: Bnd-0.0.223
Bnd-LastModified: 1198893954164
Bundle-Version: 0
```

The `Private-Package` header does appear here, because bnd copies it, but it will be ignored by the OSGi framework. One of the other headers inserted by bnd is very important though: the `Import-Package` header.


Just as packages are not made available from bundles except when explicitly exported using the `Export-Package` header, bundles cannot use the classes from a package unless they have explicitly imported that package using the `Import-Package` header (or the `Require-Bundle` header, which we will see shortly). Since the bundle we are building currently has code-level dependency on the package `org.osgi.book.reader.api`, it must import that package.

All packages that we use in the code for a bundle *must* be imported, except for packages beginning with `java.*`, which must *not* be imported. The `java.*` packages, and only those packages, are always available without being imported. There is a common misconception that packages in the standard Java runtime APIs do not need to be imported. *They do*. For example, if a bundle uses Swing, it must import `javax.swing`, but it need not import `java.awt`. If a bundle uses SAX parsing for XML documents, it must import `org.xml.sax`.

If we represent an exported package diagrammatically as follows:



and an imported package as follows:



then the runtime resolution of the two bundles will be as in Figure 3.1. The thick, dashed line in this figure represents the “wiring” together by the framework of the import with its corresponding export.

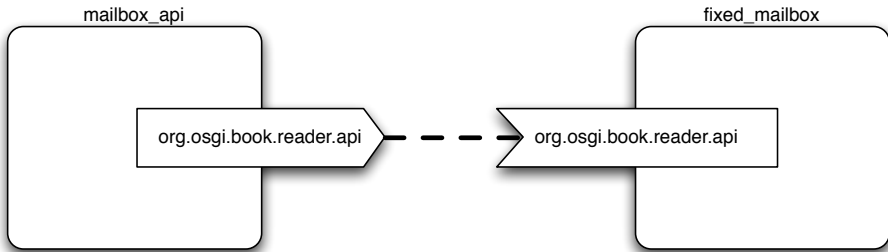


Figure 3.1: The runtime resolution of matching import and export.

A package import is one of the *constraints* that the framework must satisfy in order to resolve a bundle, i.e. move it from the `INSTALLED` state to the `RESOLVED` state. In this case the constraint on the `fixed_mailbox` bundle was satisfied by matching it with the corresponding export on the `mailbox_api` bundle. In general the framework must match all of the imported packages against packages exported from `RESOLVED` bundles before it can satisfy the import constraints of a bundle. Note that the framework can move two or more bundles into `RESOLVED` state simultaneously, making it possible to resolve bundles that have circular dependencies.

3.5 Interlude: How Bnd Works

In the previous example, the bnd descriptor for the mailbox implementation bundle contained only a single line — the `Private-Package` instruction — yet the generated manifest contained a correct `Import-Package` statement. How did that happen?

In fact this is the main features of bnd: it is able to calculate the imports of a bundle through careful inspection of the class files that we tell it to include in the bundle. Since we are making so much use of bnd, we need to take a closer look at how it generates bundles and manifests.

Bnd has several modes in which it can be used, but we will mostly be using it in its “build” mode for constructing bundles. In build mode, bnd follows a two-stage process: first it works out what the contents of the bundle should be, i.e. which packages and classes need to be included; and second, it calculates the dependencies of the included classes and generates an `Import-Package` statement.

The `Export-Package` and `Private-Package` instructions determine the contents of the bundle. Any package named in either of these instruction, either explicitly or through a wildcard pattern, will be included in the bundle JAR. Listing 3.8 shows an example.

Listing 3.8 Bnd Sample: Controlling Bundle Contents

```
# bnd sample
Export-Package: org.osgi.book.reader*
Private-Package: org.osgi.tutorial
```

This will result in all packages beginning with `org.osgi.book.reader` being included in the bundle and exported; and the package `org.osgi.tutorial` included but not exported.

Once bnd knows what should be in the JAR, it can calculate the imports. As mentioned, it inspects the bytecode of each class file found in the included packages, which yields a list of classes and packages. By default, every package found is added to the `Import-Package` header, but we can assert more control over this process by including an `Import-Package` *instruction* in our bnd file. For example, Listing 3.9 shows how to apply a specific version range to some imports and mark others as optional.

Listing 3.9 Bnd Sample: Controlling Imports

```
# bnd sample
Import-Package: org.apache.log4j*;version="[1.2.0,1.3.0)",\
               javax.swing*;resolution:=optional,\
               *
```

The entries in this list are *patterns*. Each package dependency found through bytecode analysis is tested against the patterns here in order, and if a match is found then the additional attributes are applied to the import. In this sample, if a dependency on a Log4J package is found then it will be marked with the version attribute specified; if a dependency on any Swing package is found then it will be marked as optional. The final “*” is a catch-all; any packages matching neither of the first two patterns will pass straight into the manifest without any additional attributes.

The use of wildcards like this can be initially alarming if one does not understand that the way bnd checks the actual discovered dependencies against the patterns in the `Import-Package` statement. For example the pattern `javax.swing*` may appear to be importing the whole of Swing; i.e. all 17 packages of it. If that were the case then the simple `*` on its own would be even more alarming! The confusion arises because bnd’s instructions, such as

`Import-Package`, have exactly the same name as OSGi manifest headers, but they are treated quite differently.

We can also explicitly add and remove packages from the imports:

```
# bnd sample
Import-Package: !org.foo,\
               com.bar,\
               *
```


This results in `org.foo` being excluded from the imports list, even if `bnd` detects a dependency on it! This is dangerous, but useful if the dependency exists only in a part of the code that we know can never be reached. The second entry inserts the package `com.bar` into the imports list, whether or not `bnd` is able to detect a dependency on it. This is mainly useful in cases where the bundle code uses reflection to dynamically load classes. Note that only explicitly named classes can be used: if we wrote `com.bar*`, that would be a pattern, and only included if a package starting with `com.bar` were to be found in bytecode analysis.

3.6 Requiring a Bundle

Another kind of constraint that can be placed on a bundle is the `Require-Bundle` header. This is similar to `Import-Package` header in that it makes exported packages from another bundle available to our bundle, but it works on the whole bundle rather than individual packages:

```
Require-Bundle: mailbox-api
```

If we represent a required bundle as follows:



Then the runtime resolution of a bundle using `Required-Bundle` looks like Figure 3.2. In this figure, Bundle B has a `Require-Bundle` dependency on Bundle A, meaning that Bundle B cannot resolve unless Bundle A is in `RESOLVED` state.

The effect at runtime is as if Bundle B had declared an `Import-Package` header naming every package exported by Bundle A. However, Bundle B is giving up a lot of control, because the list of imports is determined by the set of packages exported by Bundle A. If additional exports are added to Bundle A, then they are automatically added as imports to Bundle B.

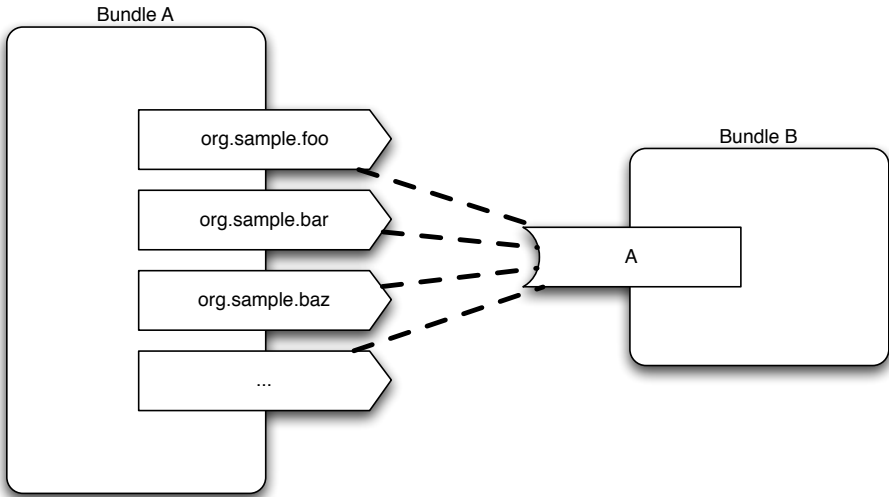


Figure 3.2: The runtime resolution of a Required Bundle

The use of `Require-Bundle` is strongly discouraged by most OSGi practitioners except where absolutely necessary, because there are several flaws with this kind of dependency scheme.

First, a bundle using `Require-Bundle` to import code from another bundle is at the mercy of what is provided by that bundle — something that can change over time. We really have no idea what packages will be provided by another bundle at any point in the future, yet nevertheless our bundle will successfully resolve even if the required bundle stops exporting some functionality that we rely on. The result will almost certainly be class loading errors such as `ClassNotFoundException` or `NoClassDefFoundError` arising in bundles that were previously working.

The second and closely related problem is that requiring bundles limits our ability to refactor the composition of those bundles. Suppose at some point we notice that Bundle *A* has grown too big, and some of the functionality it provides is really unrelated to the core and should be separated into a new bundle, which we will call Bundle *A'*. As a result, some of the exports of *A* move into *A'*. For any consumers of that functionality who are using purely `Import-Package`, the refactoring of *A* into *A* and *A'* will be entirely transparent: the imports will simply be wired differently at runtime by the framework. However, consumers requiring Bundle *A* will break, since they will no longer get the exports which have moved to *A'*.

Third, whole-module dependency systems tend to cause a high degree of “fan-out”. When a bundle requires another bundle, we have no idea (except by low-

level examination of the code) which part of the required bundle it is actually using. This can result in us bringing in a large amount of functionality when only a small amount is really required. The required bundle may also require several other bundles, and those bundles require yet more bundles, and so on. In this way we can easily be required to pull in fifty or more bundles just to resolve a single, small bundle. Using purely `Import-Package` gives us the opportunity to break this fan-out by finding instances where only a small portion of a large bundle is used, and either splitting that bundle or finding an alternative supplier for the functionality we need.

Essentially, using `Require-Bundle` is like grabbing hold of the wrapper around what we need, rather than the contents itself. It might actually work if JARs in Java were more cohesive, and the constituent packages did not change over time, but that is not the case.

`Require-Bundle` was only recently introduced into OSGi in Release 4, and given all the problems listed, it may seem mysterious that it was introduced at all. The main reason was to support various legacy issues in Eclipse, which abandoned an earlier module system in favour of OSGi. The pre-OSGi module system used by Eclipse was based on whole-module dependencies, and if OSGi had offered only `Import-Package` then the challenges of making thousands of existing Eclipse plug-ins work as OSGi bundles would have been insurmountable. Nevertheless, the existence of `Require-Bundle` remains highly controversial, and there are very, very few good reasons ever to use it in new bundles.

3.7 Version Numbers and Ranges

Versioning is a critically important part of OSGi. Modules, or libraries, are not immortal and unchanging entities — they evolve as their requirements change and as new features are added.

To take the example of Apache Log4J again, the most widely used version is 1.2. Earlier versions than this are obsolete, and there were also significant changes in the API between 1.1 and 1.2, so any code that has been written to use version 1.2 will almost certainly not work on 1.1 or earlier. There is also a 1.3 version, but it introduced a number of incompatibilities with 1.2 and is now discontinued — it is considered a failed experiment, and never produced a “stable” release. Finally there is a new version 2.0, but it is still experimental and not widely used.

This illustrates that we cannot simply use a bundle without caring about which version we wish to use. Therefore OSGi provides features which allow us first to describe in a consistent way the versions of all our bundles and their

exports, and second to allow bundles to describe the range of versions that are acceptable for each of their dependencies.

3.7.1 Version Numbers

OSGi follows a consistent scheme for all version numbers. It uses three numeric segments plus one alphanumeric segment, where any segment may be omitted. For example the following are all valid versions in OSGi:

- 1
- 1.2
- 1.2.3
- 1.2.3.beta_3

The three numeric segments are known as the *major*, *minor* and *micro* numbers and the final alphanumeric segment is known as the *qualifier*. When any one of the numeric segments is missing, it takes the implicit value of zero, so 1 is equivalent to 1.0 and 1.0.0. When a version string is not supplied at all, the version 0.0.0 is implied.

Versions have a total ordering, using a simple algorithm which descends from the major version to the qualifier. The first segment with a difference in value between two versions “short-circuits” the comparison, so later segments need not be compared. For example 2.0.0 is considered higher than 1.999.999, and this is decided without even looking at the minor or micro levels.

Things get interesting when we consider the qualifier, which may contain letters A-Z in upper or lower case, numbers, hyphens (-) and underscores (_). The qualifier is compared lexicographically using the algorithm found in the `compareTo()` method of the standard Java String class. There are two points to beware of: qualifier strings are compared character by character until a difference is found, and shorter strings are considered lower in value than longer strings.

Suppose you are getting ready to release version 2.0 of a library. This is a significant new version and you want to get it right, so you go through a series of “alpha” and “beta” releases. The alphas are numbered as follows:

- 2.0.0.alpha1
- 2.0.0.alpha2
- ...

This works fine up until the ninth alpha, but then when you release version 2.0.0.alpha10 you find that it doesn’t appear to be the highest version! This is because the number ten starts with the digit 1, which comes before the digit

2. So version `alpha10` will actually come between versions `alpha1` and `alpha2`. Therefore, if you need to use a number component inside the qualifier, always be sure to include some leading zeros. Assuming that 99 alpha releases are enough, we should have started with `alpha01`, `alpha02` and so on.

Finally after lots of testing, you are ready to unleash the final release version, which you call simply `2.0.0`. Sadly this doesn't work either, as it comes before *all* of the alpha and beta releases. The qualifier is now the empty string, which comes before all non-empty strings. So we need to add a qualifier such as `final` to the final release version.

Another approach that can work is to always add a date-based qualifier, which can be generated as part of the build process. The date would need to be written “backwards” — i.e. year number first, then month, then day and perhaps time — to ensure the lexicographic ordering matches the temporal order. For example:

- `2.0.0.2008-04-28_1230`

This approach scales well, so it is especially useful for projects that release frequently. For example, Eclipse follows a very similar scheme to this. However this approach can be inconvenient because the version strings are so verbose, and it's difficult to tell which versions are important releases versus mere development snapshots.

3.7.2 Versioning Bundles

A version number can be given to a bundle by supplying the `Bundle-Version` manifest header:

```
1 Bundle-Version: 1.2.3.alpha
```

Bundle-level versioning is important because of `Require-Bundle`, and also because OSGi allows multiple versions of the same bundle to be present in the framework simultaneously. A bundle can be uniquely identified by the combination of its `Bundle-SymbolicName` and its `Bundle-Version`.

3.7.3 Versioning Packages

However, bundle-level versioning is not enough. The recommended way to describe dependencies is with `Import-Package`, and a single bundle could contain implementations of multiple different APIs. Therefore we need version information at the package level as well. OSGi allows us to do this by tagging each exported package with a version attribute, as follows:

```
1 Export-Package: org.osgi.book.reader.api;version="1.2.3.alpha",
2   org.osgi.book.reader.util;version="1.2.3.alpha"
```

Unfortunately in the manifest file, the version attributes must be added to each exported package individually. However `bnd` gives us a shortcut:

```
# bnd sample
Export-Package: org.osgi.book.reader*;version="1.2.3.alpha"
```

We can also, if we wish to, keep the bundle version and the package export versions synchronized:

```
# bnd sample
ver: 1.2.3.alpha
Bundle-Version: ${ver}
Export-Package: org.osgi.book.reader*;version=${ver}
```

3.7.4 Version Ranges

When we import a package or require a bundle, it would be too restrictive to only target a single specific version. Instead we need to specify a range.

Recall the discussion of Apache Log4J and its various versions. Suppose we wish to depend on the stable release, version 1.2. However, “1.2” is not just a single version, it is in fact a range from 1.2.0 through to 1.2.15, the latest at the time of writing, meaning there have been fifteen “point” releases since the main 1.2 release. However, as in most projects, Log4J tries to avoid changes in the API between point releases, instead confining itself to bug fixes and perhaps minor API tweaks that will not break backwards compatibility for clients. Therefore it is likely that our code will work with any of the 1.2 point releases, so we describe our dependency on 1.2 using a version *range*.

A range is expressed as a floor and a ceiling, enclosed on each side either by a bracket “[” or a parenthesis “(”. A bracket indicates that the range is inclusive of the floor or ceiling value, whereas a parenthesis indicates it is exclusive. For example:

- [1.0.0,2.0.0)

This range includes the value 1.0.0, because of the opening bracket. but excludes the value 2.0.0 because of the closing parenthesis. Informally we could write it as “1.x”.

If we write a single version number where a range is expected, the framework still interprets this as a range but with a ceiling of infinity. In other words 1.0.0 would be interpreted as the range [1.0.0,∞). To specify a single exact version, we have to write it twice, as follows: [1.2.3,1.2.3].

For reference, here are some further examples of ranges, and what they mean when compared to an arbitrary version *x*:

[1.2.3,4.5.6)	$1.2.3 \leq x < 4.5.6$
[1.2.3,4.5.6]	$1.2.3 \leq x \leq 4.5.6$
(1.2.3,4.5.6)	$1.2.3 < x < 4.5.6$
(1.2.3,4.5.6]	$1.2.3 < x \leq 4.5.6$
1.2.3	$1.2.3 \leq x$
	$0.0.0 \leq x$

In our Log4J example and in real world usage, the first of these styles is most useful. By specifying the range [1.2,1.3) we can match any of the 1.2.x point releases. However, this may be a leap of faith: such a range would match all future versions in the 1.2 series, e.g. version 1.2.999 (if it were ever written) and beyond. We cannot test against all future versions of a library, so we must simply trust the developers of the library not to introduce breaking changes into a future point release. If we don't or can't trust those developers, then we must specify a range which includes only the known versions, such as [1.2,1.2.15].

In general it is probably best to go with the open-ended version range in most cases. The cost in terms of lost flexibility with the more conservative closed range outweighs the risk of breaking changes in future versions.

3.7.5 Versioning Import-Package and Require-Bundle

We can add version ranges to the `Import-Package` statement in exactly the same way as we added a version to `Export-Package`, using an attribute:

```
Import-Package: org.apache.log4j;version="[1.2,1.3)",
               org.apache.log4j.config;version="[1.2,1.3)"
```

Again, `bnd` can help to reduce the verbosity:

```
# bnd
Import-Package: org.apache.log4j*;version="[1.2,1.3)"
```

We can also add a version range when requiring a bundle, but the attribute name is slightly different:

```
Require-Bundle: mailbox-api;bundle-version="[1.0.0,1.1.0)"
```

It's always possible that a package import could match against two or more exports from different bundles, or a required bundle could match two or more bundles. In that case, the framework chooses the provider with the highest version, so long as that version is in the range specified by the consumer. Unfortunately, even this rule sometimes fails to come up with a single winner, because two bundles can export the same version of a package. When that happens the framework arbitrarily chooses the one with the lowest bundle ID, which tends to map to whichever was installed first.

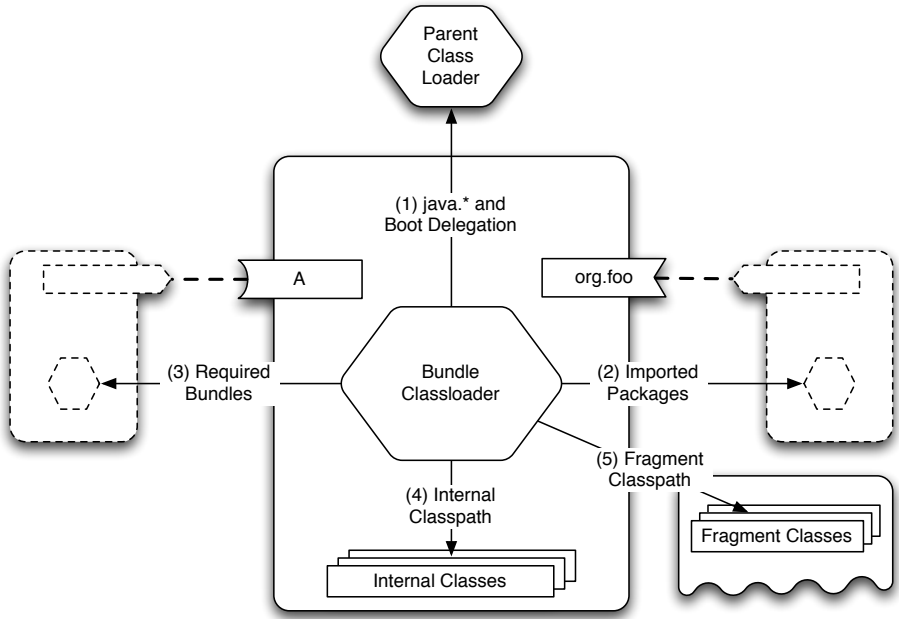


Figure 3.3: Simplified OSGi Class Search Order

3.8 Class Loading in OSGi

In Section 1.2.1 in the Introduction, we saw how normal hierarchical class loading works in Java. Figure 3.3 shows, in slightly simplified form, the process by which classes are searched in OSGi.

All resolved bundles have a class loader, and that when that class loader needs to load a class it first checks whether the package name begins with `java.*` or is listed in a special configuration property, `org.osgi.framework.bootdelegation`. If that’s the case, then the bundle class loader immediately delegates to its parent class loader, which is usually the “application” class loader, familiar from traditional Java class loading.

Why does OSGi include this element of hierarchical class loading, if it is supposed to be based on a graph of dependencies? There are two reasons:

- The only class loader permitted to define classes for the `java.*` packages is the system bootstrap class loader. This rule is enforced by the Java Virtual Machine, and if any other class loader (e.g. an OSGi bundle class loader) attempts to define one of these classes it will receive a `SecurityException`.

- Unfortunately some Java Virtual Machines, including most versions of Sun’s VM and OpenJDK, rely on the incorrect assumption that parent delegation always occurs. Because of this assumption, some internal VM classes expect to be able to find certain other internal classes through any arbitrary class loader. Therefore OSGi provides the `org.osgi.framework.bootdelegation` property to allow for parent delegation to occur for a limited set of packages, i.e. those providing the internal VM classes.

The second step that the bundle class loader takes when searching for a class is to check whether it is in a package imported with `Import-Package`. If so then it “follows the wire” to the bundle exporting the package and delegates the class loading request to that bundle’s class loader. The exporting bundle’s class loader will run the same procedure, and as a result may itself delegate to another bundle’s class loader. This means it is quite valid for a bundle to re-export a package that it itself imports.

The third step is to check whether the class is in a package imported using `Require-Bundle`. If it is then it the class loading request is delegated to the required bundle’s class loader.

Fourth, the class loader checks the bundle’s own internal classes, i.e. the classes inside its JAR. It might seem unusual that the classes that are present in the bundle JAR itself are loaded so late in the process, and can therefore be overridden by imports from other bundles. In fact this is the analogue of the “parent first” delegation that happens in traditional hierarchical class loading, and it helps to ensure class space consistency.

Fifth, the class loader searches the internal classes of any *fragments* that might be currently attached to the bundle. Fragments will be discussed in Section 3.11.

There is a sixth step that is omitted by Figure 3.3, which is related to dynamic class loading. This is an advanced topic that will be discussed in Chapter ??.

Figure 3.4 shows the full class search order in the form of a flowchart. This diagram derived from Figure 3.18 in the OSGi R4.1 Core Specification.

The class search algorithm as described always attempts to load classes from another bundle in preference to classes that may be on the classpath of the present bundle. This may seem counterintuitive at first, but in fact it makes a lot of sense, and it fits perfectly with the conventional Java approach. A traditional Java class loader always first delegates to its parent before attempting to define a class itself. This helps to ensure that classes are loaded as few times as possible, by favouring already-defined copies of a class over redefining it in a lower-level class loader. The opposite approach would result in many copies of the same class, which would be considered incompatible because they were defined by different class loaders. Back in OSGi, where the class loaders are arranged in a graph rather than a tree, the same principle of minimising

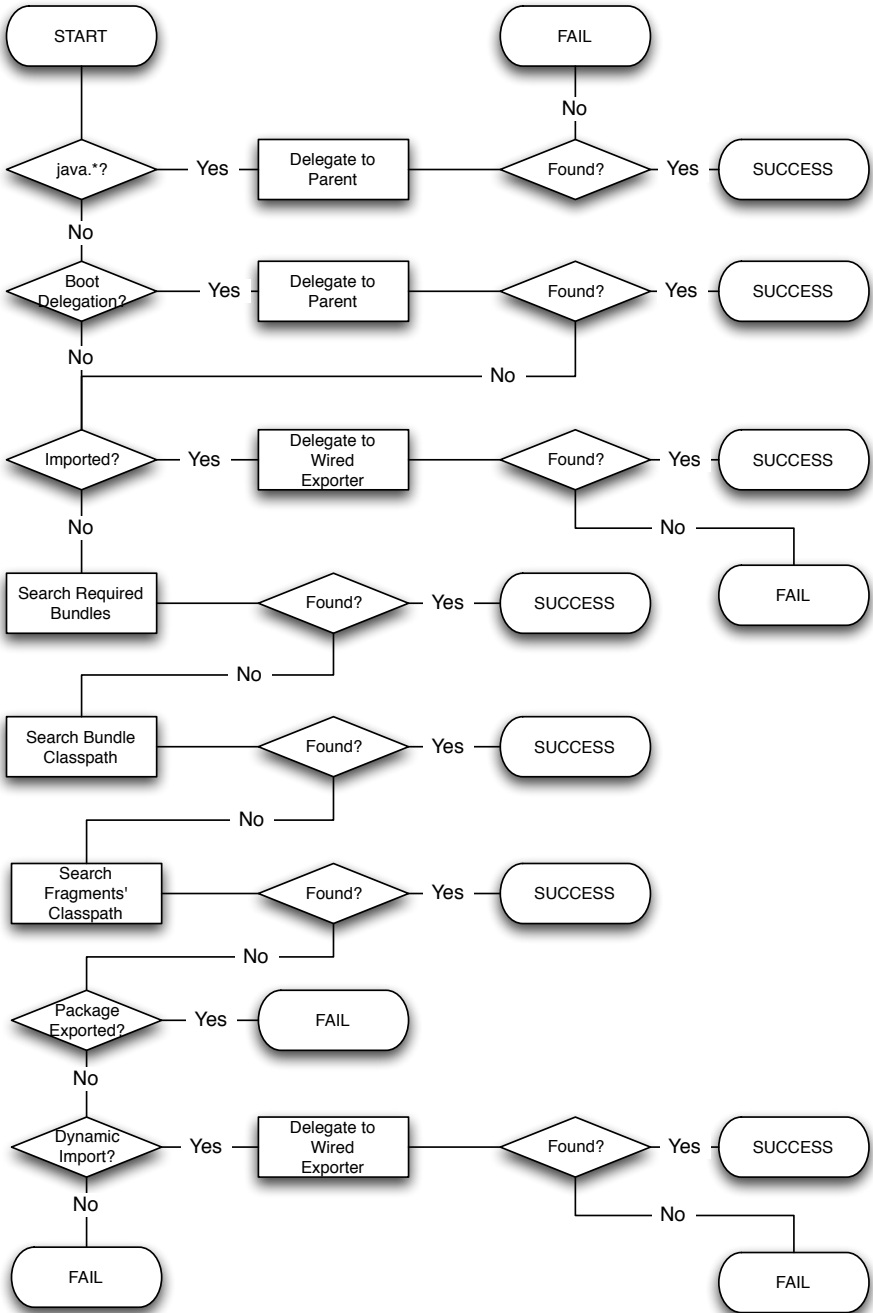


Figure 3.4: Full OSGi Search Order

class loading translates to making every effort to find a class from a bundle's imports rather than loading it from the internal bundle classpath.

Remember, the identity of a class is defined by the combination of its fully qualified name *and* the class loader which loaded it, so class `org.foo.Bar` loaded by class loader *A* is considered different from `org.foo.Bar` loaded by class loader *B*, even if they were loaded from the same physical bytes on disk. When this happens the result can be very confusing, for example a `ClassCastException` being thrown on assignment of a value of type `org.foo.Bar` to a variable of type `org.foo.Bar`!

3.9 JRE Packages

In Section 3.4 we saw that it is necessary to always import all packages used in the bundle except for `java.*`. Therefore all of the other packages in the base JRE libraries, e.g. those beginning `javax.*`, `org.omg.*`, `org.w3c.*` etc must all be imported if used.

As ordinary imports, these packages are not subject to parent delegation, so they must be supplied in the normal way by wiring the import to a bundle that provides them as an export. However, you do not need to create that bundle yourself: the system bundle performs this task.

Recall from Section 2.5 that the system bundle is always present, and it represents the framework itself as a bundle. One of the jobs performed by the system bundle is to export packages from the base JRE libraries. As it is a special bundle, the class loader inside the system bundle follows different rules and is able to load from the main application or bootstrap class loaders.

However, the list of packages exported by the system bundle is not fixed, but subject to configuration: if you take a look at Felix's `config.properties` file you will see a property named `org.osgi.framework.system.packages` which lists the packages to be exported from the system bundle. The list is different depending on which version and edition of Java we are running, because later versions of Java have more packages in their API.

This approach of importing JRE libraries using the normal wiring system provides a very clean approach to environmental dependencies. An advantage is that the importing bundle doesn't necessarily know or care that the exporter of a package is the system bundle. Some APIs which are distributed as part of the base JRE libraries in recent versions of Java can still be used on older Java versions if they are obtained as separate JAR files. For example, Java 6 includes scripting capabilities² in the `javax.script` package, however there is also a JAR available³ which implements the same functionality and works

²as defined by JSR 223[?]

³From <https://scripting.dev.java.net/>

on Java 1.4 and Java 5. We could take that JAR and create a bundle from it, so that if another bundle wishes to use scripting functionality it can simply import the `javax.script` package.

3.10 Execution Environments

Unfortunately not all environmental dependencies can be handled through importing packages. As well as the `javax.*` packages, different versions of Java have different sets of `java.*` packages, different classes in those packages and even different fields and methods in those classes.

Some of these changes are easy to catch. For example, Java 5 introduced several new methods on core classes such as `java.lang.String`. The following code will not compile on Java 1.4 because the `contains()` method was added in Java 5:

```
public class Java5Test {
    public boolean stringTest(String s) {
        return s.contains("foo");
    }
}
```

Java provides a simple defence mechanism against us compiling this code with a Java 5 compiler and then running it on a Java 1.4 environment, in the form of the class file version. The Java 5 compiler will tag all class files it creates with version 49.0, which will be rejected by a Java 1.4 VM as it expects the version to be no higher than 48.0.

In principle we could work around this mechanism using the `-source` and `-target` flags of the `javac` compiler, in which case the class would successfully load in the 1.4 VM but throw a `NoSuchMethodError` when it reached the `contains()` method. However it's difficult to imagine why anybody would inflict this kind of pain on themselves.

Unfortunately this defence mechanism still kicks in rather later than we would like it to. In OSGi terms, our bundle will resolve without any issues, but as soon as the offending class is loaded — which could happen at any arbitrary time after bundle resolution — we will get an exception. Even the error message is unhelpful:

```
Exception in thread "main" java.lang.UnsupportedClassVersionError:
Java5Test (Unsupported major.minor version 49.0)
    at java.lang.ClassLoader.defineClass0(Native Method)
    at java.lang.ClassLoader.defineClass(ClassLoader.java:539)
    at ....
```

It would be far better if we could see a clear error message “this code requires Java 5” and if that error happened at bundle resolution time.

Another problem is that a simple class version number cannot accurately represent the complex evolution tree of Java. There's not only the line of standard edition releases from 1.1 through to 6, but also the mobile edition releases CDC 1.0, CDC 1.1 and various profiles such as Foundation 1.0 and 1.1, PersonalJava 1.1 and 1.2, and so on. Even the linear evolution of Java SE may begin to look more complex in the future if Google's Android[?] or open source forks of OpenJDK pick up significant market share. Ideally, we would like to be able to specify exactly which of these environments our bundle runs on.

OSGi offers this functionality in the form of another kind of bundle resolution constraint. We can specify a list of the "execution environments" supported by our bundles, and at runtime the current Java VM must either be one of those environments or strictly upwards compatible with one of them, otherwise the bundle will fail to resolve. Furthermore OSGi tools will tell us clearly why the bundle cannot be resolved. In return we must ensure that our bundles only use packages, classes and methods that are available in *all* of the listed environments.

There is no fixed list of execution environments, since it is subject to change as the JCP creates new versions and editions of Java. However, the following set is currently supported by all three open source OSGi implementations:

- CDC-1.0/Foundation-1.0
- CDC-1.1/Foundation-1.1
- JRE-1.1
- J2SE-1.2
- J2SE-1.3
- J2SE-1.4
- J2SE-1.5
- J2SE-1.6
- OSGi/Minimum-1.0
- OSGi/Minimum-1.1

For example to specify that our bundle requires Java 5, we can add the following line to the bnd descriptor:

```
Bundle-RequiredExecutionEnvironment: J2SE-1.5
```

This entry will simply be copied verbatim to the `MANIFEST.MF` and recognised by the framework.

The first eight environments should need little explanation — they simply correspond directly with major releases of the Java specification. However the last two, OSGi/Minimum-1.0 and 1.1, are artificial environments that represent a subset of both Java SE and CDC 1.0 or 1.1. By choosing to

target the Minimum-1.0 environment we can make our bundle run essentially everywhere, or at least everywhere that OSGi itself can run.

It's a good idea to add an execution environment declaration to every bundle you build. However this is only half the problem. Having declared that our bundle requires a particular environment we need to have processes in place to ensure our code really does only use features from that environment.

It is a common misconception that the `-source` and `-target` flags of the compiler can be used to produce, say, Java 1.4-compatible code using a Java 5 compiler. They cannot. The `-source` flag controls only language features, so setting it to 1.4 turns off generics, annotations, for-each loops, and so on. The `-target` flag controls the class file version, so setting this to 1.4 makes the class file readable by a 1.4 VM. Neither of these flags do anything to restrict the APIs used from code, so we can still quite easily produce code that calls `String.contains()` or any other Java 5-only APIs.

The *only* practical and reliable way to produce 1.4-compatible code is to build it with version 1.4 of the JDK. Likewise for 1.3-compatible code, 5-compatible code, etc. The key is the `rt.jar` which contains the JRE library for that version of the JDK: only by compiling with the correct `rt.jar` in the classpath can we ensure API compatibility with the desired version. Since `rt.jar` is not available as a separate download, we must obtain the whole JDK.

3.11 Fragment Bundles

In Release 4, OSGi introduced the concept of Fragment bundles. A fragment, as its name suggests, is a kind of incomplete bundle. It cannot do anything on its own: it must attach to a host bundle, which must itself be a full, non-fragment bundle. When attached, the fragment can add classes or resources to the host bundle. At runtime, its classes are merged into the internal classpath of the host bundle.

What are fragments useful for? Here is one possibility: imagine you have a library which requires some platform-specific code. That is, the implementation of the library differs slightly when on Windows versus Mac OS X or Linux. A good example of this is the SWT library for building graphical user interfaces, which uses the “native” widget toolkit supplied by the operating system.

We don't want to create entirely separate bundles for each platform, because much of the code is platform independent, and thus would have to be duplicated. So what about isolating the platform independent code into its own bundle (let's call this bundle *A*), and making separate bundles for the platform-specific portion (which we will call P_{Win} , P_{MacOS} , etc.)? Unfortunately this separation is not always easy to achieve, because the platform-specific bundles

P_* may have lots of dependencies on the platform independent bundle, A , and in order to import those dependencies into P_* , they would have to be exported by A . However, those exports are really internal features of the library; we don't want any other bundles except P_* to access them. But once exported, they are available to anybody — in OSGi, there is no way to export packages only to specific importers.

The solution to this problem is to make P_* into fragments hosted by A rather than fully fledged bundles. Suppose we are running on Windows. In this case, the fragment P_{Win} will be merged at runtime into the internal classath of A . Now P_{Win} has full visibility of all packages in A , *including* non-exported packages, and likewise A has full visibility of all packages in P_{Win} .

Another use case is providing resources that differ depending on the locale of the user. Typically GUI applications need to provide resource bundles — usually properties files — containing all of the natural-language strings in the GUI. By separating these resources into fragments of the host bundle, we can save disk space and download time by delivering only the language fragment that the user wants.

We can even deliver different functionality for different locales. For example, written sentences in Japanese contain no spaces, yet still have discrete words which must not be split by a line break. The challenge for a word processor is to know where it is valid to insert a line break: special heuristic algorithms must be used. English may seem simpler, but even here we need to work out where long words can sensibly be split by a hyphen. Using fragments, we can separate this functionality from the base language-independent word processing functionality, and deliver only the most appropriate set of fragments for the language the user wishes to use.

3.12 Class Space Consistency and “Uses” Constraints

TODO

4 Introduction to Services

OSGi provides one of the strongest module systems in any language or platform, and we saw just a small amount of its power in the previous section. However what we have seen so far only addresses the management of static, compiled-in dependencies between modules.

This alone is not enough. A module system that only supports static dependencies is fragile and fails to offer extensibility. Fragility means that we cannot remove any single module without breaking almost the entire graph. Lack of extensibility means we cannot add to the functionality of a system without recompiling parts of the existing system to make it aware of new components.

To build a module system that is both robust and extensible, we need the ability to perform *late binding*.

4.1 Late Binding in Java

One of the most important goals of the Object Oriented Programming movement is to increase the flexibility and reuse of code by reducing the coupling between the providers of functionality (objects) and the consumers of that functionality (clients).

In Java, we approach this goal through the use of *interfaces*. An interface is a purely abstract class which provides no functionality itself, but allows any class to implement it by providing a defined list of methods. By coding against interfaces rather than concrete types, we can write code that isolates itself almost entirely from any specific implementation.

For example, suppose we wish to write a mailbox scanner component, which periodically scans a mailbox for new messages. If we design our `Mailbox` interface carefully then we can write a scanner that neither knows nor cares which specific type of mailbox it is scanning — whether it be an IMAP mailbox, an RSS/ATOM “mailbox”, an SMS mailbox, etc. Therefore we can reuse the same scanner component for all of these mailbox types without changing its code at all. Because of these benefits, so-called “interface based” programming is now widely recognised as good Java programming practice.

However, there is a catch: we cannot create new objects unless we know their specific type. Interfaces and abstract classes can be used to refer to an object

after it is created, but they cannot be used to instantiate new ones, because we need to know exactly what type of thing to create. This would cause a problem if our scanner component expects to create the mailbox object itself. A naïve solution might be to put some kind of switch in the constructor of the scanner to make it decide at runtime what kind of mailbox to create, as shown in Listing 4.1.

Listing 4.1 Naïve Solution to Instantiating an Interface

```

1 public class MailboxScanner {
2     private final Mailbox mailbox;
3     public MailboxScanner(String mailboxType) {
4         if("imap".equals(mailboxType)) {
5             mailbox = new IMAPMailbox();
6         } else if("rss".equals(mailboxType)) {
7             mailbox = new RSSMailbox();
8         } else {
9             // ...

```

Most programmers would recognise this as a terrible idea, and in OSGi it's even worse, because it will only work if the bundle containing the scanner imports *every* package that might contain a mailbox implementation class... including ones that might be written in the future!

The normal solution to this quandary is for the scanner not to try to instantiate the mailbox itself, but to allow a mailbox to be supplied to it by something else. This is the essence of late binding: the consumer of functionality is not bound to a specific provider until runtime. But who or what should this “something else” be?

There are many possible answers. A large application may have hundreds of classes like `MailboxScanner` which all need to be supplied with implementations of the interfaces they depend on, so a common theme in many solutions is to centralise object creation into an “Assembler” class. That central class may even support some form of scripting, to allow the network of objects to be rewired without recompilation.

4.1.1 Dependency Injection Frameworks

After writing a few examples of the “Assembler” class in different projects, it's easy to see that it is a common pattern that can be extracted out into a framework. And indeed several open source frameworks have emerged that do exactly this: popular examples in Java are the Spring Framework[?] and Google Guice[?].

Such a framework is often called a “Dependency Injection” (DI) framework¹. Unfortunately both these and the manually-created Assembler pattern have

¹They have also previously been called “Inversion of Control” (IoC) frameworks, but this

traditionally suffered from being mostly static: the wiring together of “beans” (i.e., plain Java objects, such as the mailbox implementation) with their consumers (other beans, such as the mailbox scanner) tends to happen once, at start-up of the application, and remains in force until the application is shut-down.

Static wiring results in many problems. The first is fragility due to a sensitive dependence on start-up ordering. For example if object *B* depends on object *A*, then *A* must be created before *B*. When we scale up to thousands of objects and their interdependencies, the dependency graph becomes brittle and far too complex for any human to understand. We must avoid circular dependencies between objects, at all costs.

Another problem is the impossibility of on-the-fly updates. Most production systems need to be patched occasionally as bugs are fixed or requirements change. In a statically wired dependency graph, we usually need to shut down the entire system even when updating only the tiniest corner of it.

4.1.2 Dynamic Services

OSGi solves these problems with dynamic *services*.

A service, like a bean in a DI framework, is a plain Java object. It is published in the OSGi Service Registry under the name of one or more Java interfaces, and consumers who wish to use it may look it up using any of those interfaces names. Services may consume other services, but rather than being wired into a fixed graph, services can be registered or unregistered dynamically at any time, so they form only temporary associations with each other.

Start ordering problems can now be solved easily: services start in any order. Suppose we start the bundle containing service *B* before starting the bundle containing service *A*. In this case, *B* simply waits for *A* to become available. Also we can update individual components without restarting the system. When taking away service *A* and replacing it with *A'*, OSGi sends events to service *B* to keep it informed of the situation.

Services offer an interface-based programming model. The only requirement on a service is that it implements an interface; *any* interface will do, even ones from the base JRE or third-party libraries. The chosen interface — or interfaces, since a Java object can implement many interfaces — forms the primary addressing mechanism for a service. For example, a service publisher declares “I have this object available which is `Mailbox`.” The consumers declare “I am looking for a `Mailbox`.” The Service Registry provides a venue for

term has largely fallen out of use thanks to Martin Fowler, who popularised the term “Dependency Injection” in his 2004 article “Inversion of Control Containers and the Dependency Injection pattern”[?].




Figure 4.1: Service Oriented Architecture

publishers and consumers to find each other. The consumer does not need to know the implementation class of the published service, just the interface through which it interacts.

OSGi's Service Registry has been called a form of Service Oriented Architecture, or SOA. Many people think of SOA as being associated with distributed computing, Web Services, SOAP and so on, but that is just one example of SOA, which is really just a pattern or style of architecture. OSGi services are limited in scope to a single JVM — they are not distributed² — yet they map very cleanly to the concepts of SOA. Figure 4.1 is adapted from a diagram used by the World-Wide Web Consortium to explain SOA, and we will see shortly how each of the abstract concepts in that diagram maps to a concrete entity in OSGi.

Sadly, there is some complexity cost involved in handling dynamic services versus static objects. It is obviously easier to make use of something if we know that it is always available! However, the real world is not static, and therefore we need to write systems that are robust enough to handle entities that come and go. The good news is that several high-level abstractions have been built on top of the OSGi Service Registry that greatly simplify the code you need to write while still taking advantage of dynamic behaviour. In fact, the Dependency Injection frameworks have started to support dynamic behaviour, and in the case of the Spring Framework that support is achieved through direct usage of the OSGi Service Registry.

²Although of course some people have sought to build distributed systems on top of OSGi services

In this chapter though, we look at the nuts and bolts of services, in order to gain a firm foundation when advancing to higher-level abstractions.

4.2 Registering a Service

Recall that in Section 3.4 we created an implementation of the `Mailbox` interface. However, as its package was not exported, there was no way for any other bundle to actually use that implementation! By registering an instance of the `FixedMailbox` class as a service, other bundles can access the object without needing to have a dependency on its class.

Registering a service is an example of interaction with the OSGi framework, and as always we need to have a `BundleContext` in order to do that. Therefore we need to write an implementation of `BundleActivator` as shown in Listing 4.2.

Listing 4.2 Welcome Mailbox Activator

```
1 package org.osgi.book.reader.fixedmailbox;
2
3 import org.osgi.book.reader.api.Mailbox;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6
7 public class WelcomeMailboxActivator implements BundleActivator {
8
9     public void start(BundleContext context) throws Exception {
10         Mailbox mbox = new FixedMailbox();
11         context.registerService(Mailbox.class.getName(), mbox, null); //1
12     }
13
14     public void stop(BundleContext context) throws Exception {
15     }
16 }
```

The first two lines of code simply create a mailbox with some hard-coded messages in it. The line marked `//1` is the one that calls the framework to register the mailbox as a service, by calling the `registerService` method with three parameters:

1. The *interface name* under which the service is to be registered. This should be the name of a Java interface class, and will be the primary means by which clients will find the service.
2. The *service object* itself, i.e. the mailbox object that was instantiated in the previous lines. This object *must* implement the interface named in the previous parameter.
3. A set of *service properties*, which here has been left blank by passing `null`. We will look at service properties shortly.

We say that a service object is registered “under” an interface name, because for consumers the most important fact about our service is the interface it implements. In fact, just as Java object can be implementations of multiple interfaces, we can register a service under multiple interface names, by calling a variant of the `registerService` method that takes an array of `Strings` as its first parameter. When we do this, consumers can find the service using any one of those interface names. There is still just one service object registered: the additional entries can be considered aliases.

Note that we could have passed a literal `String` for the interface name i.e., `context.registerService("org.osgi.book.reader.api.Mailbox", mbox, null)`. However this is not good practice, because the class or package name might change in the future, for example when the code is refactored. It is better to use `Mailbox.class.getName()` because most IDEs can automatically update the import statement — and even if we are not using an IDE, we will get a helpful compilation error if we change the package name that `Mailbox` lives in without updating this source file.

Listing 4.3 Bnd Descriptor for the Welcome Mailbox Bundle

```
# welcome_mailbox.bnd
Private-Package: org.osgi.book.reader.fixedmailbox
Bundle-Activator: \
    org.osgi.book.reader.fixedmailbox.WelcomeMailboxActivator
```

Let’s build this bundle and see what effect it has. The `bnd` descriptor should look like the one in Listing 4.3. After building, installing and starting the bundle, try typing “`services 11`” where “11” should be substituted with the actual bundle ID of the bundle (use the `ps` command to find it). You should see this result:

```
-> services 11

welcome_mailbox (11) provides:

objectClass = org.osgi.book.reader.api.Mailbox
service.id = 24
```

This shows we have successfully registered a service under the `Mailbox` interface, and it has been assigned a service ID of 24 — again, you will probably get a different ID when you run this for yourself. Incidentally you can try the `services` command on its own, without a bundle ID parameter, which will give you a list of all registered services by all bundles, although showing less detail.

4.3 Unregistering a Service

In the code listing in Section 4.2, we did not do anything in the `stop()` method of the activator, which may seem oddly asymmetrical. Usually we have to undo in the `stop()` method whatever we did in the `start()` method, so shouldn't we have to unregister the service?

Actually this is not necessary because the OSGi framework automatically unregisters any services registered by our bundle when it is deactivated. We don't need to explicitly clean it up ourselves, so long as we are happy for the service's lifecycle — the period of time during which it was registered and available to clients — to coincide with the lifecycle of the bundle itself. Sometimes though we need a different lifecycle for the service. We may only want to offer the service when some other conditions are met, and in that case we will have to control both registering and unregistering ourselves.

Suppose, for example, we only want our service to be available while a particular file exists on the filesystem. Perhaps that file contains some messages which we want to offer as a mailbox: clearly we can only offer the service if the file actually exists. To achieve this we would create a polling thread using the same pattern we saw in Section 2.11. The code is shown in Listing 4.4.

Here we see, at marker *1*, that the `registerService()` method returns an object of type `ServiceRegistration`, and we can use that object to unregister the service later. Each pass through the loop we check whether the file `messages.txt` exists in your home directory³. If it does, and the service is not currently registered, then we register it. If the file does *not* exist, and the service *is* currently registered, then we unregister it.

For the purposes of this example, the actual implementation of `FileMailbox` is not particularly interesting. If you simply wish to get this code working and watch the service register and unregister, then you will need to create a simple stub implementation. The code in Listing 4.5 will suffice, assuming we never actually call the methods of the `FileMailbox` object.

The `bind` descriptor should be as in Listing 4.6.

By creating and deleting the `messages.txt` file, you should be able to see the service appear and disappear, albeit with up to five seconds' delay. Incidentally if we were to provide a real implementation of `FileMailbox`, this delay would be one of the things we would have to take into account in order to make the service robust: we should be able to handle the situation where a client request arrives during the period between the file being deleted and the service being

³On Windows XP, this is usually `C:\Documents and Settings\YourName` and on Windows Vista it is `C:\Users\YourName`. Users of non-Windows operating systems tend to know how to find their home directory already.

Listing 4.4 File Mailbox Activator

```
1 package org.osgi.book.reader.filemailbox;
2
3 import java.io.File;
4 import java.util.Properties;
5
6 import org.osgi.book.reader.api.Mailbox;
7 import org.osgi.framework.BundleActivator;
8 import org.osgi.framework.BundleContext;
9 import org.osgi.framework.ServiceRegistration;
10
11 public class FileMailboxActivator implements BundleActivator {
12
13     private Thread thread;
14
15     public void start(BundleContext context) throws Exception {
16         File file = new File(System.getProperty("user.home")
17             + System.getProperty("file.separator") + "messages.txt");
18         RegistrationRunnable runnable = new RegistrationRunnable(
19             context, file, null);
20         thread = new Thread(runnable);
21         thread.start();
22     }
23
24     public void stop(BundleContext context) throws Exception {
25         thread.interrupt();
26     }
27 }
28
29 class RegistrationRunnable implements Runnable {
30
31     private final BundleContext context;
32     private final File file;
33     private final Properties props;
34
35     public RegistrationRunnable(BundleContext context, File file,
36         Properties props) {
37         this.context = context;
38         this.file = file;
39         this.props = props;
40     }
41
42     public void run() {
43         ServiceRegistration registration = null;
44         try {
45             while (!Thread.currentThread().isInterrupted()) {
46                 if (file.exists()) {
47                     if (registration == null) {
48                         registration = context.registerService( // 1
49                             Mailbox.class.getName(),
50                             new FileMailbox(file), props);
51                     }
52                 } else {
53                     if (registration != null) {
54                         registration.unregister();
55                         registration = null;
56                     }
57                 }
58                 Thread.sleep(5000);
59             }
60         } catch (InterruptedException e) {
61             // Allow thread to exit
62         }
63     }
64 }
```

Listing 4.5 File Mailbox (Stub Implementation)

```
1 package org.osgi.book.reader.filemailbox;
2
3 import java.io.File;
4
5 import org.osgi.book.reader.api.Mailbox;
6 import org.osgi.book.reader.api.Message;
7
8 /**
9  * Warning: Empty stub implementation
10 */
11 public class FileMailbox implements Mailbox {
12
13     private static final long[] EMPTY = new long[0];
14
15     public FileMailbox(File file) {}
16     public long[] getAllMessages() { return EMPTY; }
17     public Message[] getMessages(long[] ids) {
18         return new Message[0];
19     }
20     public long[] getMessagesSince(long id) { return EMPTY; }
21     public void markRead(boolean read, long[] ids) { }
22 }
```

Listing 4.6 Bnd Descriptor for File Mailbox Bundle

```
# file_mailbox.bnd
Private-Package: org.osgi.book.reader.filemailbox
Bundle-Activator: org.osgi.book.reader.filemailbox.FileMailboxActivator
```

unregistered. In that case we would have to return an error message to the client.

4.4 Looking up a Service

Having seen how to register and unregister services, the next logical step is to look at how to look up and call methods on those services.

Perhaps surprising, this can be a little tricky. The problem is that services, as we have seen, can come and go at any time. If we look for a service at a particular instant, we might not find it, but we would find it if we looked two seconds later. Alternatively we could access a service twice in a row, but get a different service the second time because the one we got first time is no longer around.

Fortunately there is lots of support, both in the OSGi specifications themselves and in external third-party libraries, to help us to abstract away this complexity. In fact programming with dynamic services in OSGi need be hardly any more complex than with static dependency injection, yet it is far more powerful. However, in this section we will look at the most low-level way of accessing services. This is partially to provide a firm base of understanding for when we get onto the more convenient approaches, and partially to drive home the truly dynamic nature of OSGi services.

Suppose we wish to write a bundle that accesses one of the `Mailbox` services and prints the current total number of messages in that mailbox. To keep things as simple as possible we will do this in the `start()` method of an activator, for example `MessageCountActivator` in Listing 4.7.

At marker *1*, we ask the framework to find a `ServiceReference` for a named Java interface. As before, we avoid encoding the interface name as a literal string.

After checking that the service reference is not `null`, meaning the mailbox service is not currently available, we proceed to request the actual service object at marker *2*. Again we have to check if the framework returned `null` — this is because although the service was available at marker *1*, it may have become unavailable in the time it took us to reach marker *2*.

At marker *3* we actually call the service. Because our `mbox` variable holds the actual service object rather than any kind of proxy, we can call the methods on it just like any normal Java object.

Finally at marker *4* we “un-get” the service — in other words we let the framework know that we are no longer using it. This is necessary because the framework maintains a count of how many bundles are using a particular

Listing 4.7 Message Count Activator

```
1 package org.osgi.tutorial;
2
3 import org.osgi.book.reader.api.Mailbox;
4 import org.osgi.book.reader.api.MailboxException;
5 import org.osgi.framework.BundleActivator;
6 import org.osgi.framework.BundleContext;
7 import org.osgi.framework.ServiceReference;
8
9 public class MessageCountActivator implements BundleActivator {
10
11     private BundleContext context;
12
13     public void start(BundleContext context) throws Exception {
14         this.context = context;
15         printMessageCount();
16     }
17
18     public void stop(BundleContext context) throws Exception {
19     }
20
21     private void printMessageCount() throws MailboxException {
22         ServiceReference ref = context // 1
23             .getServiceReference(Mailbox.class.getName());
24
25         if (ref != null) {
26             Mailbox mbox = (Mailbox) context.getService(ref); // 2
27             if (mbox != null) {
28                 try {
29                     int count = mbox.getAllMessages().length; // 3
30                     System.out.println("There are " + count + "messages");
31                 } finally {
32                     context.ungetService(ref); // 4
33                 }
34             }
35         }
36     }
37 }
```

service: when that count is zero, it knows that the service can safely be removed when the bundle providing it is deactivated. We have placed the call to `ungetService()` inside a `finally` block to ensure that it is always called on exit from our method, even if an uncaught exception occurs.

Listing 4.8 shows the Bnd descriptor.

Listing 4.8 Bnd Descriptor for the Message Counter Bundle

```
# message_count.bnd
Private-Package: org.osgi.tutorial
Bundle-Activator: org.osgi.tutorial.MessageCountActivator
```

Stepping back from the code, you may wonder why accessing a service requires a two-stage process of first obtaining a reference and then obtaining the actual service object. We will look at the reasons for this in the next section.

The main problem with this code is that it's just too long! Simply to make a single call to a service, we have to write two separate `null` checks and a `try / finally` block, along with several noisy method calls to the OSGi framework. We certainly don't want to repeat all of this each time we access a service. Also, the code does not behave particularly well when the `Mailbox` service is unavailable: it simply gives up and prints nothing. We could at least print an error message, but even that is unsatisfactory: what if we really *need* to know how many messages are in the mailbox? The information will be available as soon as the mailbox service is registered, but it's just bad luck that the above bundle activator has been called first.

4.5 Service Properties

In addition to the interface name, we may wish to associate additional metadata with our services. For example, in this chapter we have registered several instances of the `Mailbox` service, each with different characteristics. It would be nice to tell clients something about each mailbox so that they can, if they wish, obtain only one specific mailbox, or at least report something to the user about what kind of mailbox they have obtained. This is done with service properties.

Recall that when registering the service in Section 4.2, we left the final parameter `null`. Instead we can pass in a `java.util.Properties` object containing the properties that we want to set on the service⁴.

⁴Actually the type of this parameter is `java.util.Dictionary`, of which `java.util.Properties` is a sub-class. Why not use the `java.util.Map` interface, which the `Properties` class also implements? Simply because `Map` has “only” existed since Java 1.2, so it cannot be used on all platforms supported by OSGi!

Let's modify `WelcomeMailboxActivator` to add a "mailbox name" property to the service. The new `start` method is shown in Listing 4.9.

Listing 4.9 Adding Service Properties to the Welcome Mailbox

```
11 public void start(BundleContext context) throws Exception {
12     Mailbox mbox = new FixedMailbox();
13
14     Properties props = new Properties();
15     props.put(Mailbox.NAME_PROPERTY, "welcome");
16     context.registerService(Mailbox.class.getName(), mbox, props);
17 }
```

Note that we avoid hard-coding the property name everywhere we use it, as this can be error prone and difficult to change later. Instead we use a constant that was defined in the `Mailbox` interface itself. This is a suitable place to put the constant because it is guaranteed to be visible to both service implementers and service consumers.

Try rebuilding this bundle and updating it in Felix. If we now enter the command "`services 11`" we should see the property has been added to the service:

```
-> services 11

welcome_mailbox (11) provides:
-----
objectClass = org.osgi.book.reader.api.Mailbox
service.id = 24
mailboxName = welcome
```

The other entries we see here, `objectClass` and `service.id` are built-in properties that have been added by the framework. There are other standard property names defined in the OSGi specification, but these are the only two compulsory ones, so they appear on every service. These two, and the rest, can be found in the class `org.osgi.framework.Constants`.

To query the properties on a service, we simply call the `getProperty()` method of the `ServiceReference` to get a single property value. If we need to know about all of the properties on the service we can call `getPropertyKeys()` method to list them.

Service properties provide a clue as to why we need a two-stage process for looking up a service, i.e. first obtaining a `ServiceReference` before obtaining the actual service. One reason for this is sometimes we don't need the service object (or at least not *yet*) but only its properties. When we obtain the service object, the OSGi framework must keep track of the fact we are using it, creating a very small amount of overhead. It is sensible to avoid that overhead when it is not needed.

Also, service references are small, lightweight objects that can be passed around, copied or discarded at will. The framework does not track the service

reference objects it hands out. Also service reference objects can be passed easily between bundles, whereas it can cause problems to pass an actual service instance to another bundle: that would prevent the framework from being able to track which bundles are using each service. In situations where we want to ask another bundle to do something with a service, we should pass the reference and let the other bundle call `getService()` and `ungetService()` itself.

4.6 Introduction to Service Trackers

To address the excess verbosity of the code in Section 4.4, we can refactor it to use a utility class provided by OSGi called `ServiceTracker`

In fact `ServiceTracker` is one of the most important classes you will use as an OSGi programmer, and it has many more uses than just the one we will be taking advantage of in this code. But for now, the code in Listing 4.10 simply does the same thing as before, and in particular it is no better at dealing with a missing `Mailbox` service:

Listing 4.10 Message Count Activator — `ServiceTracker` version

```

1 package org.osgi.tutorial;

3 import org.osgi.book.reader.api.Mailbox;
4 import org.osgi.book.reader.api.MailboxException;
5 import org.osgi.framework.BundleActivator;
6 import org.osgi.framework.BundleContext;
7 import org.osgi.util.tracker.ServiceTracker;

9 public class MessageCountActivator2 implements BundleActivator {

11     private ServiceTracker mboxTracker;

13     public void start(BundleContext context) throws Exception {
14         mboxTracker = new ServiceTracker(context, Mailbox.class // 1
15             .getName(), null);
16         mboxTracker.open(); // 2
17         printMessageCount();
18     }

20     public void stop(BundleContext context) throws Exception { // 3
21         mboxTracker.close();
22     }

24     private void printMessageCount() throws MailboxException {
25         Mailbox mbox = (Mailbox) mboxTracker.getService(); // 4
26         if (mbox != null) {
27             int count = mbox.getAllMessages().length; // 5
28             System.out.println("There are " + count + " messages");
29         }
30     }
31 }

```

At first glance this code is barely any shorter than the previous example!

Nevertheless we have achieved something important: in exchange for a little extra initial effort, it is now much easier to call the service, as we can see in the method `printMessageCount`. In real code, we would probably make many separate calls to the service, so it is far better to repeat the four lines of code in this `printMessageCount` than the nine lines of code in the previous version.

The first difference, which we can see at marker 1 of the `start` method, is that instead of saving the bundle context directly into a field, we instead construct a new `ServiceTracker` field, passing it the bundle context and the name of the service that we are using it to track. Next at marker 2, we “open” the tracker, and at marker 3 in the `stop()` method we “close” it. We will look later at what is really going on under the covers when we open and close a service tracker, but for now just remember that the tracker will not work until it is opened.

The next difference is at marker 4, where we call `getService` on the tracker. Refreshingly, this immediately gives us the actual service object (if available) rather than a `ServiceReference`. So we simply go ahead and call the service at marker 5, bearing in mind that we still need to check if the service was found. Also, we don’t need to clean up after ourselves with a `finally` block: we simply let the variable go out of scope, as the tracker will take care of releasing the service.

Here’s another interesting thing we can do with `ServiceTracker`. In the version of `printMessageCount` shown in Listing 4.11 we don’t want to fail immediately if the service is unavailable; instead we would like to wait up to five seconds for the service to become available.

Listing 4.11 Waiting for a Service

```
36 private void printMessageCount(String message)
37     throws InterruptedException, MailboxException {
38     Mailbox mbox = (Mailbox) mboxTracker.waitForService(5000);
39     if (mbox != null) {
40         int count = mbox.getAllMessages().length;
41         System.out.println("There are " + count + " messages");
42     }
43 }
```

When the log service is available, `waitForService` works exactly the same as `getService`: it will immediately return the service instance. However if the log service is not currently available, the call will block until either the service becomes available or 5000 milliseconds has passed, whichever is sooner. Only after 5000 milliseconds has passed without the service becoming available will it return `null`.

However, `waitForService` needs to be used with caution, and in particular it should not be called from the `start` or `stop` methods of a `BundleActivator`,

because those methods are supposed to return control quickly to the framework. This topic is discussed in more depth in Section ??.

4.7 Listening to Services

In both versions of the “message counting” bundle above, we failed to handle a missing mailbox service gracefully. If the mailbox wasn’t there, we simply gave up.

Sometimes giving up is exactly the right thing to do. For example there is a standard log service in OSGi which clients can use to send logging messages to an application-wide log. Suppose a component wants to write a message to the log, but it discovers that the log service is not currently available — what should it do? Usually it should give up, i.e. not write the message to the log, but carry on with its main task. Logs are nice if we can have them, but should not get in the way of actually running the application.

However at other times, we need to do better. Some components are essentially useless without the services that they consume: a component that counts messages in a mailbox is pointless when there are no mailboxes. We might say that it has a *dependency* on the mailbox service.

Services often depend on other services. For example, suppose we have a relational database that is exposed to our application as a service under the `javax.sql.DataSource` interface — we might then wish to offer a mailbox service that obtains its messages from the database. In that case, the “database mailbox” service would have a dependency on the `DataSource` service, and it would be useless when the `DataSource` is not available. This is very similar to the `FileMailbox` that was discussed in Section 4.3, in which we registered the service only when the backing file was available, and unregistered the service when it was not. In the case of the database mailbox, we can register the service when the `DataSource` service is available, and unregister when it is not.

In the file example, we had to write a thread to poll for the existence of the backing file. Polling is not ideal because it wastes CPU time, and there is inevitably a delay between the state of the file changing and our program detecting that change. But there is currently no way in Java — short of using platform-specific native code — to “watch” or “listen” to the state of a file⁵, so polling is our only choice. Not so with services. By listening to the service registry, we can be notified immediately when the `DataSource` service is registered or unregistered, and react immediately by registering or unregistering the corresponding mailbox.

⁵JSR 203[?] (also known as “NIO.2”) seeks to address this limitation for Java 7.

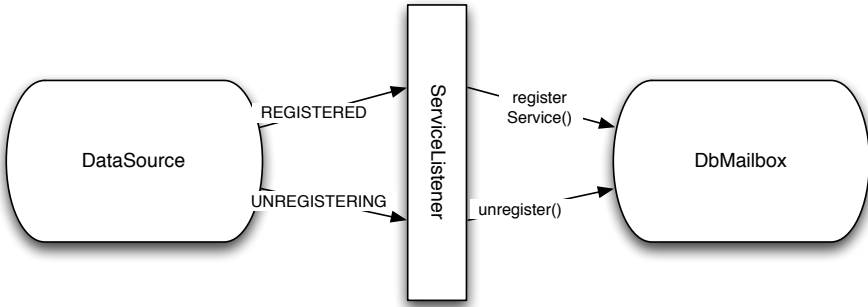


Figure 4.2: Updating a Service Registration in Response to Another Service

To achieve this, we listen to events published by the service registry. Whenever a service is either registered or unregistered, the framework publishes a `ServiceEvent` to all registered `ServiceListeners`. Any bundle can register a `ServiceListener` through its `BundleContext`. Therefore we could write a simple listener that, when it receives an event of type `REGISTERING`, registers a new `DbMailbox` service, and when it receives an event of type `UNREGISTERING`, unregisters that service.

Unfortunately, this will not work.

The problem is that service listeners are only told about state *changes*, not about pre-existing state. If the `DataSource` is already registered as a service before we start our listener, then we will never receive the `REGISTERING` event, but nevertheless we need to register the `DbMailbox` service. Therefore we really need to do two things:

- Scan the service registry for a pre-existing service under the `DataSource` interface; if such exists, then register a `DbMailbox` for it.
- Hook up a service listener which will register a `DbMailbox` when a `DataSource` service is registered, and unregister it when the `DataSource` service is unregistered

But if we do those two things in the stated order, it will *still* not work! There will be a window of time between scanning the pre-existing services and starting the service listener, and if a service is registered during that window, we will miss it completely. Therefore we need to reverse the two steps: *first* hook up the listener, *then* scan the pre-existing services. Now there is the potential for overlap: if a service registers between the two steps it will be handled both by the scanning code and also by the listener, so we need to be careful to guard against duplication.

All this sounds nightmarishly complex, and indeed the code needed to get all of this right is very unpleasant. We are not even going to look at an example

because you should never have to write any code like this. Instead you should be using `ServiceTracker`.

4.8 Tracking Services

Although we saw `ServiceTracker` in Section 4.6, we only used it in a very limited way. In this section we will see the main purpose for `ServiceTracker`: hiding the complexities of listening to and consuming dynamic services. Rather than simply “listening”, which is passive, we wish to actively “track” the services we depend on.

Before proceeding to the example code, we will need to have an implementation of the `DbMailbox` class but, like `FileMailbox`, the actual implementation is not interesting as part of this exposition. Therefore we will use another stub class — the full definition of which is left as an exercise.

Let’s write a bundle activator using `ServiceTracker`: this is shown in Listing 4.11. The `start` and `stop` methods of this activator look very similar to our first example of using a tracker in Section 4.6: we simply create and open the tracker on start-up (markers 1 and 2), and close it on shutdown (marker 3). However this time the third parameter to the constructor of `ServiceTracker` is not `null`. Instead, we pass in an instance of the `ServiceTrackerCustomizer` interface, which tells the tracker what to do when services are added, removed or modified.

Our customizer simply registers a `DbMailbox` service whenever a `DataSource` service is added, and unregisters it when the `DataSource` service is removed.

Why does this not suffer from the same limitations as the `ServiceListener`-based approach described Section 4.7? Simply because, unlike a listener, the adding and removed methods of `ServiceTracker` are called not only when the state of a service changes but also when the tracker is opened, to notify us of pre-existing services. The `addingService()` method is called multiple times when the tracker is opened, once for each service currently registered, and it is also called whenever a new service is registered at any time later for as long as the tracker is open. Furthermore the `removedService()` is called any time a service that we have been previously notified of goes away, and it is also called for each service when the tracker closes. Therefore we can deal with services in a uniform fashion without needing to distinguish between pre-existing services and ones that are registered while our listener is active. This greatly simplifies the code we need to write.

Incidentally the `ServiceTracker` class does not use any special hooks into the framework, it builds on the existing facilities that we have already seen. When we call `open()` on a service tracker, it hooks up a `ServiceListener` and then

Listing 4.12 Database Mailbox Activator

```
1 package org.osgi.book.reader.dbmailbox;
3 import javax.sql.DataSource;
5 import org.osgi.book.reader.api.Mailbox;
6 import org.osgi.framework.BundleActivator;
7 import org.osgi.framework.BundleContext;
8 import org.osgi.framework.ServiceReference;
9 import org.osgi.framework.ServiceRegistration;
10 import org.osgi.util.tracker.ServiceTracker;
11 import org.osgi.util.tracker.ServiceTrackerCustomizer;
13 public class DbMailboxActivator implements BundleActivator {
15     private BundleContext context;
16     private ServiceTracker tracker;
18     public void start(BundleContext context) throws Exception {
19         this.context = context;
21         tracker = new ServiceTracker(context, DataSource.class
22             .getName(), new DSCustomizer());           // 1
23         tracker.open();                                 // 2
24     }
26     public void stop(BundleContext context) throws Exception {
27         tracker.close();                               // 3
28     }
30     private class DSCustomizer implements ServiceTrackerCustomizer {
32         public Object addingService(ServiceReference ref) {
33             DataSource ds = (DataSource) context.getService(ref); // 4
35             DbMailbox mbox = new DbMailbox(ds);
36             ServiceRegistration registration = context.registerService(
37                 Mailbox.class.getName(), mbox, null);           // 5
39             return registration;                               // 6
40         }
42         public void modifiedService(ServiceReference ref,
43             Object service) {
44         }
46         public void removedService(ServiceReference ref, Object service) {
47             ServiceRegistration registration =
48                 (ServiceRegistration) service;                 // 7
49             registration.unregister();                         // 8
50             context.ungetService(ref);                       // 9
51         }
52     }
53 }
54 }
```

scans the pre-existing services, eliminates duplicates etc. The internal code is still complex, but it has been written for us (and exhaustively tested) to save us from having to do it ourselves.

Let's look at the customizer class in a little more detail. At marker *4* we receive a `ServiceReference` from the tracker that points at the newly registered service, and we ask the framework to de-reference it to produce the actual service object. We use the result to create a new instance of `DbMailbox` and at marker *5* we register it as a new mailbox service. At marker *6* we return the registration object back to the tracker.

Why return the registration object? The signature of `addingService()` simply has a return type of `Object`; the tracker does not care what type of object we return, and we can return anything we like, including `null`. However the tracker promises to remember the object and give it back to us in the following situations:

- When we call `getService`, the tracker will return whatever object we returned from `addingService`.
- When the underlying service is modified or unregistered, the tracker will call `modifiedService` or `removedService` respectively, passing both the service reference and the object that we returned from `addingService`.

Because of the first one of these promises, it's perhaps more conventional to return the actual service object from `addingService` — which in the code above would be the `ds` variable. But that is not a requirement. In general we should return whatever will be most useful for us to find the information or data structure that might need to be updated later. In the above code we want to “update” the registration of the mailbox service, so we return the registration object. Other times we might be storing information into a `Map`, so we return one of the keys from the `Map`. The only exception to this rule is when we return `null`, which the tracker takes to mean that we don't care about this particular service reference. That is, if we return `null` from `addingService`, the tracker will “forget” that particular service reference and will not call either `modifiedService` or `removedService` later if it is modified or removed.

Since we know that the second parameter of `removedService` will be of type `ServiceRegistration`, we are able to cast it back to that type at marker *7*. Then we can simply unregister it (marker *8*) and “unget” the `DataSource` service (marker *9*). The final step is necessary as the mirror image of calling `getService()` in the adding method.

4.9 Filtering on Properties

Section 4.5 described how properties can be added to services when they are registered, and how the properties on a `ServiceReference` can be introspected. Now let's look at another important use for properties: filtering service look-ups.

Both of the look-up code samples we saw, in Sections 4.4 and 4.6, obtained a single instance of the mailbox service. Yet it should be clear by now that there can be an arbitrary number of service instances for each particular type, because any bundle can register a new service under that type. Therefore it is sometimes necessary to further restrict the set of services obtained by a look-up. This is done with *filters*, which are applied to the properties of the service. A filter is a simple string, using a format which is very easy to construct either manually or programmatically.

For example, suppose we wish to find the “welcome” mailbox service, and not any other kind of mailbox. Recall that that mailbox service has a property named `mailboxName` with the value “welcome”. The filter string required to find this service is simply:

```
(mailboxName=welcome)
```

Suppose we added a further property to the welcome mailbox indicating the language. To find the English version, which should have the `lang` property set to “en”, we construct the following composite filter:

```
(&(mailboxName=welcome) (lang=en))
```

Some languages have variants, such as `en_UK` and `en_US` for British and American English respectively. Suppose we want to match any kind of English:

```
(&(mailboxName=welcome) (lang=en*))
```

Finally, suppose we want either German (“de”) or any form of English *except* Canadian:

```
(&(mailboxName=welcome) (|(lang=de)(lang=en*)) (!(lang=en_CA)))
```

This syntax is borrowed directly from LDAP search filters, as defined in [?]. A filter is either a simple operation, or a composite. Here are some examples of simple operations:

<code>(foo=*)</code>	Property <code>foo</code> is present
<code>(foo=bar)</code>	Value of property <code>foo</code> is equal to “bar”
<code>(count>=1)</code>	Value of property <code>count</code> is 1 or greater
<code>(count<=10)</code>	Value of property <code>count</code> is 10 or less
<code>(foo=bar*)</code>	Value of property <code>foo</code> is a string starting “bar”...

Composite filters can be built up from simple filters, or they might compose filters which are themselves composites. Here are the composition operations:

<code>(!(filter))</code>	Boolean “NOT”: filter is false
<code>(&(filter1)...(filterN))</code>	Boolean “AND”: all filters are true
<code>((filter1)...(filterN))</code>	Boolean “OR”: any one filter is true

So how do we use these filters? It depends how we are doing the service look-up. If we are using the low-level approach from Section 4.4 then we simply use an alternative signature for the `getServiceReference()` method that takes a filter in addition to a service interface name:

```
context.getServiceReference(Mailbox.class.getName(),
    "&(mailboxName=welcome)(lang=en)");
```

If we are using a service tracker as in Sections 4.6 or 4.8, we need to use an alternative constructor for the `ServiceTracker` class which takes a `Filter` object *instead of* a service interface name. `Filter` objects can be constructed by a call to `createFilter()`:

```
Filter filter = context.createFilter(
    "&(objectClass=" + Mailbox.class.getName() + ") +
    "(mailboxName=welcome)(lang=en)");
tracker = new ServiceTracker(context, filter, null);
```

The filter object replaces the service interface name parameter because we can track instances of multiple service types with the same tracker. However when we do wish to restrict the tracker to a single service type, we need to include that constraint in the filter using the built-in property name `objectClass`. In fact, the constructor we were using previously is simply a convenience wrapper for the filter-based constructor.

While we can construct filters easily using string concatenation as above, this can be somewhat error prone — it is all too easy to miss a closing bracket. Therefore it is a good idea to build a very simple library for constructing filter strings. Listing 4.13 shows a suggestion to get started.

We can now construct our filter as shown in Listing 4.14. This code is a little longer, but it is much harder to make a mistake in the filter syntax, and it is easier to refer to constants defined elsewhere.

4.10 Cardinality and Selection Rules

Sometimes, no matter how much filtering we apply when looking up a service, we cannot avoid matching against multiple services. Since any bundle is free to register services with any set of properties, there is no way to force each service to have a unique set of properties. An exception is the service ID property, which is supplied by the framework and guaranteed to be unique,

Listing 4.13 Filter Building Utilities

```

1 package org.osgi.book.utils.filter;
2
3 public abstract class FilterBuilder {
4
5     @Override
6     public final String toString() {
7         return append(new StringBuilder()).toString();
8     }
9
10    public abstract StringBuilder append(StringBuilder builder);
11 }

```

```

1 package org.osgi.book.utils.filter;
2
3 public class EqFilter extends FilterBuilder {
4
5     private final String name;
6     private final String value;
7
8     public EqFilter(String name, String value) {
9         this.name = name;
10        this.value = value;
11    }
12    public StringBuilder append(StringBuilder buffer) {
13        return buffer.append('(').append(name).append('=')
14            .append(value).append(')');
15    }
16 }

```

```

1 package org.osgi.book.utils.filter;
2
3 public class AndFilter extends FilterBuilder {
4
5     private final FilterBuilder[] filters;
6
7     public AndFilter(FilterBuilder... filters) {
8         this.filters = filters;
9     }
10
11    public StringBuilder append(StringBuilder buffer) {
12        StringBuilder builder = new StringBuilder();
13        builder.append("&");
14        for (FilterBuilder filter : filters) {
15            filter.append(builder);
16        }
17        builder.append(" ");
18        return builder;
19    }
20 }

```

Listing 4.14 Sample Usage of Filter Builder

```

context.createFilter(new AndFilter(
    new EqFilter(Constants.OBJECTCLASS, Mailbox.class.getName()),
    new EqFilter(Mailbox.NAME_PROPERTY, "welcome"),
    new EqFilter("lang", "en*")
).toString());

```

but since we cannot know in advance what the ID of a service will be, it is not generally useful to filter on it.

Therefore the cardinality of all service look-ups in OSGi is implicitly “zero to many”. So what do we do if we prefer to simply have one?

Here’s an example: OSGi provides a standard service for sending logging messages to the system-wide log. To write messages to the log, a component can obtain an instance of the `LogService` service and call the `log` method therein. However, we would prefer to write to just one log. What do we do when there are many?

Looking back at the examples from Sections 4.4 and 4.6, it seems we don’t have to worry about this after all. The `getServiceReference` method on `BundleContext` and the `getService` method on `ServiceTracker` both return either a *single* service, or `null` if no matching services are available. They do this by applying two simple rules. The first is to look at a special `service.ranking` property on the services, which can be referenced in code as `Constants.SERVICE_RANKING`. The value of the property is an integer between `Integer.MIN_VALUE` (i.e. -2,147,483,648) and `Integer.MAX_VALUE` (2,147,483,647). The service with the highest ranking is selected — services that do not have an explicit ranking property take the implicit value of zero. If this rule produces a tie, then the service with the lowest service ID is selected. The second rule is somewhat arbitrary, but it tends to result in the “oldest” service being selected, since in most framework implementation service IDs are allocated from an incrementing counter (although this behaviour is not part of the specification and cannot be relied upon).

Concerns about cardinality are generally separated into the *minimum* number of instances required and the *maximum* instances required. If the minimum is zero then the service cardinality is optional; if the minimum is one then the service cardinality is mandatory. If the maximum is one then the cardinality is unary; if the maximum is many (i.e. there is no maximum) then the cardinality is multiple. We can refer to any of the four possible combinations as follows:

0..1	Optional and unary
0..n	Optional and multiple
1..1	Mandatory and unary
1..n	Mandatory and multiple

Let’s look at some typical usage patterns.

4.10.1 Optional, Unary

This is the simplest case: we wish to use a particular service if it is available, but don’t mind if it is not available. Also, if there are many instances, we

don't mind which one is used. A typical example of this is the standard log service mentioned above.

The normal pattern for this is the code from Section 4.6: calling `getService` against a simple (un-customized) service tracker when needed.

Another useful pattern is to sub-class `ServiceTracker` and implement the service interface directly. Listing 4.15 shows an example of this idea. Here we have an implementation of `LogService` that calls an underlying service in the registry each time a `log` method is called, if such a service exists; otherwise the log message is silently thrown away. This pattern can be useful when we want let an object use the service interface without making it aware of the details of service management. Listing 4.16 shows a simple example, where the `DatababaseConnection` class simply requires an instance of `LogService` and does not know anything about service tracking.

Listing 4.15 Log Tracker

```
1 package org.osgi.book.utils;
2
3 import org.osgi.framework.BundleContext;
4 import org.osgi.framework.ServiceReference;
5 import org.osgi.service.log.LogService;
6 import org.osgi.util.tracker.ServiceTracker;
7
8 public class LogTracker extends ServiceTracker implements LogService {
9
10     public LogTracker(BundleContext context) {
11         super(context, LogService.class.getName(), null);
12     }
13
14     public void log(int level, String message) {
15         log(null, level, message, null);
16     }
17
18     public void log(int level, String message, Throwable exception) {
19         log(null, level, message, exception);
20     }
21
22     public void log(ServiceReference sr, int level, String message) {
23         log(sr, level, message, null);
24     }
25
26     public void log(ServiceReference sr, int level, String message,
27         Throwable exception) {
28         LogService log = (LogService) getService();
29         if (log != null) {
30             log.log(sr, level, message, exception);
31         }
32     }
33 }
```

Listing 4.16 Sample Usage of Log Tracker

```
1 package org.osgi.book.utils.sample;
2
3 import org.osgi.service.log.LogService;
4
5 public class DbConnection {
6
7     private final LogService log;
8
9     public DbConnection(LogService log) {
10         this.log = log;
11         log.log(LogService.LOG_INFO, "Opening connection");
12         // ...
13     }
14
15     public void disconnect() {
16         log.log(LogService.LOG_INFO, "Disconnecting");
17         // ...
18     }
19 }
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548
2549
2550
2551
2552
2553
2554
2555
2556
2557
2558
2559
2560
2561
2562
2563
2564
2565
2566
2567
2568
2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591
2592
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607
2608
2609
2610
2611
2612
2613
2614
2615
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645
2646
2647
2648
2649
2650
2651
2652
```

4.10.2 Optional, Multiple

In this case we wish to use *all* instances of a service if any are available, but don't mind if none are available. This style is commonly used for notifications or event handlers: all of the handlers should be notified of an event, but the originator of the event doesn't care if there are no handlers. We sometimes refer to this the *Whiteboard Pattern* and we will discuss it at length in Chapter 7.

The normal implementation pattern is almost the same as for optional-unary, simply replacing the call to the `getService` method on `ServiceTracker` with a call to `getServices`. This returns an array containing all of the matching services that are currently registered. Beware that, as in other parts of the OSGi API, `getServices` returns `null` to signify no matches rather than an empty array, so we must always perform a `null`-check before using the result.

Listing 4.17 shows a variant of the `LogTracker` from the previous section. This `MultiLogTracker` will pass on logging messages to *all* available log services... as before, messages will be silently thrown away if there are no logs.

4.10.3 Mandatory, Unary

TODO

4.10.4 Mandatory, Multiple

TODO

Listing 4.17 Multi Log Tracker

```
1 package org.osgi.book.utils;

3 import org.osgi.framework.BundleContext;
4 import org.osgi.framework.ServiceReference;
5 import org.osgi.service.log.LogService;
6 import org.osgi.util.tracker.ServiceTracker;

8 public class MultiLogTracker extends ServiceTracker
9     implements LogService {

11     public MultiLogTracker(BundleContext context) {
12         super(context, LogService.class.getName(), null);
13     }

15     public void log(int level, String message) {
16         log(null, level, message, null);
17     }

19     public void log(int level, String message, Throwable exception) {
20         log(null, level, message, exception);
21     }

23     public void log(ServiceReference sr, int level, String message) {
24         log(sr, level, message, null);
25     }

27     public void log(ServiceReference sr, int level, String message,
28         Throwable exception) {
29         Object[] logs = getServices();
30         if (logs != null) {
31             for (Object log : logs) {
32                 ((LogService) log).log(sr, level, message, exception);
33             }
34         }
35     }
36 }
```

5 Example: Mailbox Reader GUI

We now have the equipment we need to embark on our task to build a GUI mailbox reader application, as first described in Section 3.1. Let's first think about the design.

The GUI should be simple and built on Java Swing. The central component will be split vertically: in the top part there will be a tabbed pane showing one tab per mailbox, and the bottom part will show the content of the selected message. Each mailbox will be displayed as a table with one row per message.

5.1 The Mailbox Table Model and Panel

We will approach the implementation of this application by building from the bottom up. First, we consider the table of messages that represents an individual mailbox. If we use the Swing `JTable` class then we need a “model” for the content of the table. The easiest way to provide this is to override the `AbstractTableModel` class and provide the few remaining abstract methods as shown in Listing 5.1. Note that this class is pure Swing code and does not illustrate any OSGi concepts, so it can be briefly skimmed if you prefer.

Given the table model class, we can now create a “panel” class that encapsulates the creation of the model, the table and a “scroll pane” to contain the table and provide it with scroll bars. This is shown in Listing 5.2.

5.2 The Mailbox Tracker

Next we will build the tracker that tracks the mailbox services and creates tables to display them. As this tracker class will be a little complex, we will build it up incrementally over the course of this section. In Listing 5.3 we simply define the class and its constructor.

Here we pass the `BundleContext` to the superclass's constructor, along with the fixed name of the service that we will be tracking, i.e. `Mailbox`. In the constructor of this tracker we also expect to receive an instance of `JTabbedPane`, which is the Swing class representing tab panels. We need this because we will

Listing 5.1 The Mailbox Table Model

```
1 package org.osgi.book.reader.gui;
2
3 import java.util.ArrayList;
4 import java.util.List;
5
6 import javax.swing.table.AbstractTableModel;
7
8 import org.osgi.book.reader.api.Mailbox;
9 import org.osgi.book.reader.api.MailboxException;
10 import org.osgi.book.reader.api.Message;
11
12 public class MailboxTableModel extends AbstractTableModel {
13
14     private static final String ERROR = "ERROR";
15
16     private final Mailbox mailbox;
17     private final List<Message> messages;
18
19     public MailboxTableModel(Mailbox mailbox) throws MailboxException {
20         this.mailbox = mailbox;
21         long[] messageIds = mailbox.getAllMessages();
22         messages = new ArrayList<Message>(messageIds.length);
23         Message[] messageArray = mailbox.getMessages(messageIds);
24         for (Message message : messageArray) {
25             messages.add(message);
26         }
27     }
28
29     public synchronized int getRowCount() {
30         return messages.size();
31     }
32
33     public int getColumnCount() {
34         return 2;
35     }
36
37     @Override
38     public String getColumnName(int column) {
39         switch (column) {
40             case 0:
41                 return "ID";
42             case 1:
43                 return "Subject";
44             }
45         return ERROR;
46     }
47
48     public synchronized Object getValueAt(int row, int column) {
49         Message message = messages.get(row);
50         switch (column) {
51             case 0:
52                 return Long.toString(message.getId());
53             case 1:
54                 return message.getSummary();
55             }
56         return ERROR;
57     }
58 }
```

Listing 5.2 Mailbox Panel

```
1 package org.osgi.book.reader.gui;
2
3 import javax.swing.JPanel;
4 import javax.swing.JScrollPane;
5 import javax.swing.JTable;
6
7 import org.osgi.book.reader.api.Mailbox;
8 import org.osgi.book.reader.api.MailboxException;
9
10 public class MailboxPanel extends JPanel {
11
12     private final MailboxTableModel tableModel;
13
14     public MailboxPanel(Mailbox mbox) throws MailboxException {
15         tableModel = new MailboxTableModel(mbox);
16         JTable table = new JTable(tableModel);
17         JScrollPane scrollPane = new JScrollPane(table);
18
19         add(scrollPane);
20     }
21 }
```

Listing 5.3 The Mailbox Tracker, Step One: Constructor

```
1 package org.osgi.book.reader.gui;
2
3 // ... import statements omitted ...
4
5 public class ScannerMailboxTracker extends ServiceTracker {
6
7     private final JTabbedPane tabbedPane;
8
9     public ScannerMailboxTracker(BundleContext ctx,
10         JTabbedPane tabbedPane) {
11         super(ctx, Mailbox.class.getName(), null);
12         this.tabbedPane = tabbedPane;
13     }
14
15     // ...
16 }
```

dynamically add tabs when each mailbox is registered and remove them when the corresponding mailbox is unregistered.

Now consider the `addingService` method. What we want to do is create a `MailboxPanel` from the mailbox, and add it to the tabbed panel. This will create a new tab. Finally, we would like to return the mailbox panel from `addingService` so that the tracker will give it back to us in `removedService`, because we will need it in order to remove the corresponding tab.

Sadly, things are not quite that simple. Like most modern GUI libraries, Swing is single-threaded, meaning we must always call its methods from a specific thread. But in OSGi, we cannot be sure which thread we are in when `addingService` is called. So instead of calling Swing directly, we must create a `Runnable` and pass it to the Swing API via `SwingUtilities.invokeLater`. This will cause the code block to be run on the correct thread, but it will happen at some unknown time in the future. Probably it will run *after* we have returned from our `addingService` method, so we cannot return the mailbox panel object directly since it may not exist yet.

One solution is to wrap the panel object in a *future*. This is a kind of “suspension”: a future represents the result of an asynchronous computation that may or may not have completed yet. In Java it is represented by the `java.util.concurrent.Future` interface, which is part of the new concurrency API introduced in Java 5. The `FutureTask` class (in the same package) is an implementation of `Future`, and it also implements `Runnable`, allowing it to be executed by a call to the Swing `invokeLater` utility method. Therefore we can write the `addingService` method as shown in Listing 5.4, returning an object of type `Future<MailboxPanel>` instead of the `MailboxPanel` directly.

Let’s examine how this method works step-by-step. The first two lines simply retrieve the mailbox object and its name. The bulk of the method constructs a `Callable` object that implements the computation we wish to perform in the GUI thread. This computation creates the mailbox panel, adds it to the tabbed panel (using the mailbox name as the tab title), and finally returns it. Returning a result from the `Callable` sets the value of the `Future`. Finally, we wrap the computation in a `FutureTask` and pass it to the Swing `invokeLater` method for execution by the GUI thread.

The `removedService` method is now quite easy: see Listing 5.5. Again we use an `invokeLater` call to update the GUI by pulling the panel out of its `Future` wrapper and removing it from the tabbed panel.

5.3 The Main Window

Now let’s look at the class that defines the main window frame, creates the tabbed panel, and uses the mailbox tracker. See Listing 5.6.

Listing 5.4 The Mailbox Tracker, Step Two: `addingService` method

```

30  @Override
31  public Object addingService(ServiceReference reference) {
32      final String mboxName =
33          (String) reference.getProperty(Mailbox.NAME_PROPERTY);
34      final Mailbox mbox = (Mailbox) context.getService(reference);

36      Callable<MailboxPanel> callable = new Callable<MailboxPanel>() {
37          public MailboxPanel call() {
38              MailboxPanel panel;
39              try {
40                  panel = new MailboxPanel(mbox);
41                  String title = (mboxName != null) ?
42                      mboxName : "<unknown>";
43                  tabbedPane.addTab(title, panel);
44              } catch (MailboxException e) {
45                  JOptionPane.showMessageDialog(tabbedPane, e.getMessage(),
46                      "Error", JOptionPane.ERROR_MESSAGE);
47                  panel = null;
48              }
49              return panel;
50          }
51      };
52      FutureTask<MailboxPanel> future =
53          new FutureTask<MailboxPanel>(callable);
54      SwingUtilities.invokeLater(future);

56      return future;
57  }

```

Listing 5.5 The Mailbox Tracker, Step Three: `removedService` method

```

59  @Override
60  public void removedService(ServiceReference reference, Object svc) {

62      @SuppressWarnings("unchecked")
63      final Future<MailboxPanel> panelRef = (Future<MailboxPanel>) svc;

65      SwingUtilities.invokeLater(new Runnable() {
66          public void run() {
67              try {
68                  MailboxPanel panel = panelRef.get();
69                  if (panel != null) {
70                      tabbedPane.remove(panel);
71                  }
72              } catch (ExecutionException e) {
73                  // The MailboxPanel was not successfully created
74              } catch (InterruptedException e) {
75                  // Restore interruption status
76                  Thread.currentThread().interrupt();
77              }
78          }
79      });

81      context.ungetService(reference);
82  }
83 }

```

Listing 5.6 The Mailbox Reader Main Window

```
1 package org.osgi.book.reader.gui;

3 import java.awt.BorderLayout;
4 import java.awt.Component;
5 import java.awt.Dimension;

7 import javax.swing.JFrame;
8 import javax.swing.JLabel;
9 import javax.swing.JTabbedPane;
10 import javax.swing.SwingConstants;

12 import org.osgi.framework.BundleContext;

14 public class ScannerFrame extends JFrame {

16     private JTabbedPane tabbedPane;
17     private ScannerMailboxTracker tracker;

19     public ScannerFrame() {
20         super("Mailbox Scanner");

22         tabbedPane = new JTabbedPane();
23         tabbedPane.addTab("Mailboxes", createIntroPanel());
24         tabbedPane.setPreferredSize(new Dimension(400, 400));

26         getContentPane().add(tabbedPane, BorderLayout.CENTER);
27     }

30     private Component createIntroPanel() {
31         JLabel label = new JLabel("Select a Mailbox");
32         label.setHorizontalAlignment(SwingConstants.CENTER);
33         return label;
34     }

37     protected void openTracking(BundleContext context) {
38         tracker = new ScannerMailboxTracker(context, tabbedPane);
39         tracker.open();
40     }

42     protected void closeTracking() {
43         tracker.close();
44     }
45 }
```

This class is also simple. In the usual way when working with Swing we subclass the `JFrame` class, which defines top-level “shell” windows. In the constructor we create the contents of the window, which at this stage is just the tabbed panel. Unfortunately tabbed panels can look a little strange when they have zero tabs, so we add a single fixed tab containing a hard-coded instruction label. In a more sophisticated version, we might wish to hide this placeholder tab when there is at least one mailbox tab showing, and re-show it if the number of mailboxes subsequently drops back to zero.

In addition to the constructor and the `createIntroPanel` utility method, we define two protected methods for managing the mailbox service tracker: `openTracking`, which creates and opens the tracker; and `closeTracking` which closes it. These methods are protected because they will only be called from the same package — specifically, they will be called by the bundle activator, which we will look at next.

5.4 The Bundle Activator

To actually execute the above code in a conventional Java Swing application, we would create a “Main” class with a `public static void main` method, as shown in Listing 5.7, and launch it by naming that class on the command line. Notice the window listener, which we need add to make sure that the Java runtime shuts down when the window is closed by the user. If this were not supplied, then the JVM would continue running even with no windows visible.

Listing 5.7 Conventional Java Approach to Launching a Swing Application

```
1 package org.osgi.book.reader.gui;
2
3 import java.awt.event.WindowAdapter;
4 import java.awt.event.WindowEvent;
5
6 public class ScannerMain {
7
8     public static void main(String[] args) {
9         ScannerFrame frame = new ScannerFrame();
10        frame.pack();
11
12        frame.addWindowListener(new WindowAdapter() {
13            public void windowClosing(WindowEvent e) {
14                System.exit(0);
15            }
16        });
17
18        frame.setVisible(true);
19    }
20
21 }
```

However in OSGi, we don't need to start and stop the whole JVM. We can use the lifecycle of a bundle to create and destroy the GUI, possibly many times during the life of a single JVM. Naturally we do this by implementing a bundle activator, and the code in Listing 5.8 demonstrates this idea.

Listing 5.8 Using Bundle Lifecycle to Launch a Swing Application

```

1 package org.osgi.book.reader.gui;

3 import java.awt.event.WindowAdapter;
4 import java.awt.event.WindowEvent;

6 import javax.swing.UIManager;

8 import org.osgi.framework.BundleActivator;
9 import org.osgi.framework.BundleContext;
10 import org.osgi.framework.BundleException;

12 public class ScannerFrameActivator implements BundleActivator {

14     private ScannerFrame frame;

16     public void start(final BundleContext context) throws Exception {
17         UIManager.setLookAndFeel(
18             UIManager.getSystemLookAndFeelClassName());
19         frame = new ScannerFrame();
20         frame.pack();

22         frame.addWindowListener(new WindowAdapter() {
23             public void windowClosing(WindowEvent e) {
24                 try {
25                     context.getBundle().stop();
26                 } catch (BundleException ei) {
27                     // Ignore
28                 }
29             }
30         });

32         frame.openTracking(context);

34         frame.setVisible(true);
35     }

37     public void stop(BundleContext context) throws Exception {
38         frame.setVisible(false);

40         frame.closeTracking();
41     }

43 }

```

When the bundle is started, we create a frame and open it by making it visible. We also tell the frame to start tracking mailbox services, passing in our bundle context. When the bundle is stopped, we hide the frame and turn off the tracker.

Notice that we again registered a window listener. This time, the listener just stops the bundle rather than stopping the entire Java runtime. The `getBundle` method on `BundleContext` returns the bundle's *own* handle object, and

bundles are free to manipulate themselves via the methods of this handle in the same way that they can manipulate other bundles. This time it is not really essential to register the listener, as it was with the standalone program, but to do so creates a pleasingly symmetry. If stopping the bundle closes the window, shouldn't closing the window stop the bundle? The lifecycles of the two are thus linked. If the user closes the window, then we can see that later by looking at the bundle state, and reopen the window by starting the bundle again.

5.5 Putting it Together

We're ready to build and run the application. Let's look at the `bnd` descriptor for the GUI application, as shown in Listing 5.9. This descriptor simple instructs `bnd` to bundle together all of the classes in the package `org.osgi.book.reader.gui` — in other words, all the classes we have seen in this chapter — and it also specifies the activator in the normal way.

Listing 5.9 Bnd Descriptor for the Mailbox Scanner

```
# mailbox_gui.bnd
Private-Package: org.osgi.book.reader.gui
Bundle-Activator: org.osgi.book.reader.gui.ScannerFrameActivator
```

We can build this bundle by running:

```
java -jar path/to/bnd.jar mailbox_gui.bnd
```

And then we can install it into Felix and start it as follows:

```
-> install mailbox_gui.jar
Bundle ID: 6
-> start 6
->
```

The `mailbox_api` bundle must also be installed, otherwise `mailbox_gui` will not resolve, and will therefore return an error when you attempt to start it. However assuming the bundle starts successfully, a Swing window should open, which looks like Figure 5.1.

Note that there are no mailboxes displayed; this is because there are no `Mailbox` services present. To view a mailbox we can install the `fixed_mailbox` bundle from Chapter 4:

```
-> install file:build/bundles/fixed_mailbox.jar
Bundle ID: 7
-> start 7
->
```

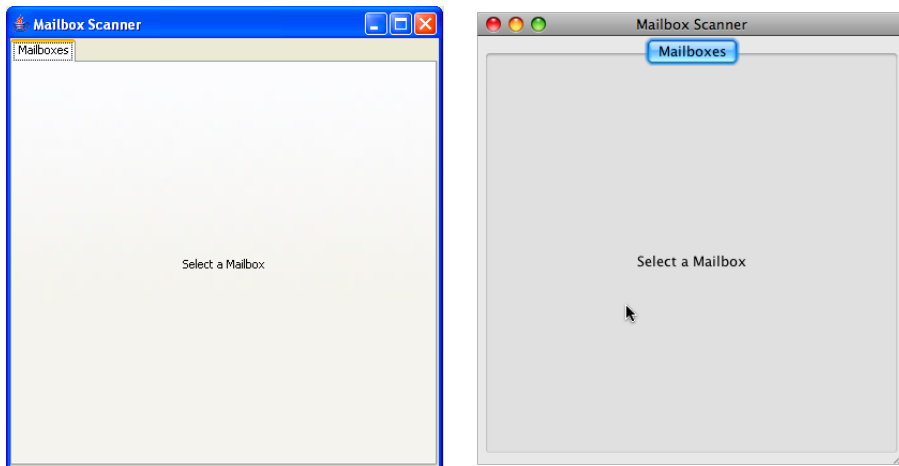


Figure 5.1: The Mailbox GUI (Windows XP and Mac OS X)

Now a tab should appear with the label “welcome”. If we select that tab, the list of messages in the mailbox will be displayed as in Figure 5.2 (from now on, only the Windows XP screenshot will be shown).

It’s worth stopping at this point to check that the application behaves as we expect it to. Specifically:

1. If we stop the `fixed_mailbox` bundle, the “welcome” tab will disappear; and if we start it again, the tab will come back.
2. If we stop the `mailbox_gui` bundle, the window will close. Starting again will re-open the window with the same set of tabs as before.
3. If we close the window, the `mailbox_gui` bundle will return to RESOLVED state (use the `ps` command to check). Again, restarting the bundle will open the window in its prior state.

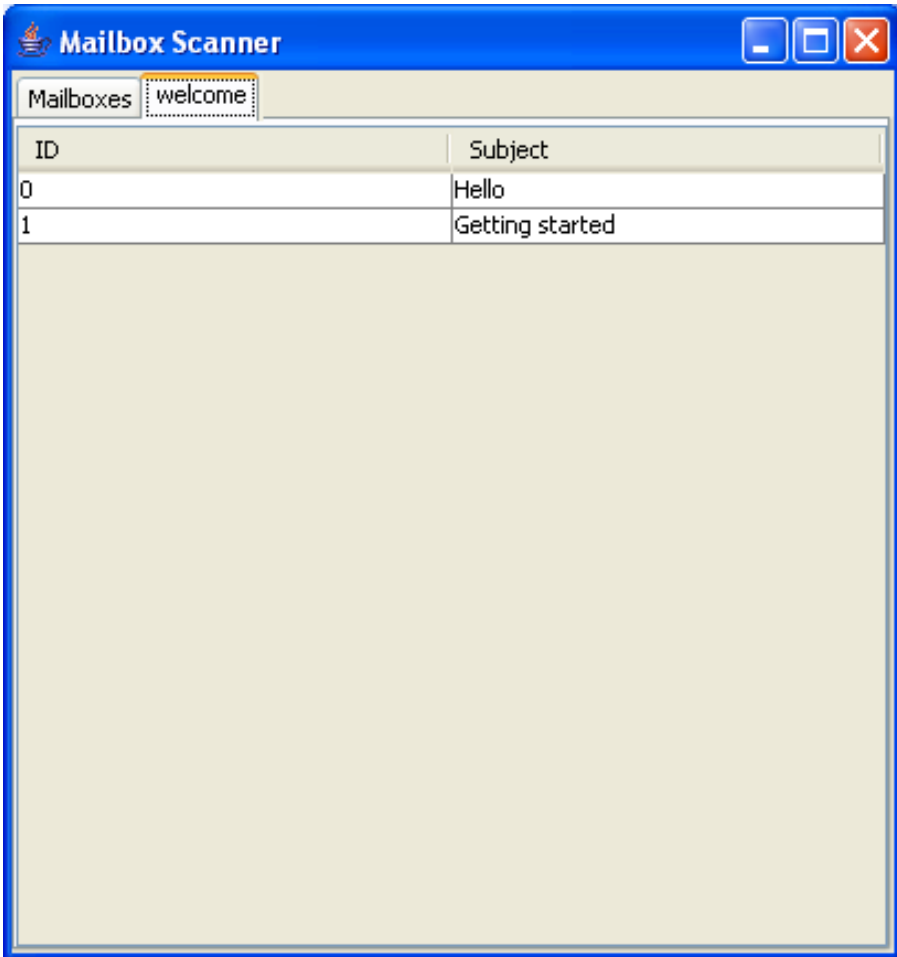


Figure 5.2: The Mailbox GUI with a Mailbox Selected

6 Concurrency and OSGi

We have touched lightly on some of the issues involved with concurrency and multi-threading in OSGi. Now it's time to look at them seriously.

Unlike heavyweight frameworks such as J2EE, OSGi does not attempt to take control of all the resources of the Java Virtual Machine, and that includes threads: whereas J2EE forbids you from writing code that creates threads or uses explicit synchronization, offering instead a cumbersome and limited “work management” framework, OSGi simply allows you to control the creation and scheduling of threads in your application yourself, as you would in any ordinary application. To support this the OSGi libraries are thread safe and can be called from any thread.

However, this freedom comes at a price. Just as we are free to create threads in our bundles, any other bundle has that freedom also. We can never assume our bundle lives in a single-threaded environment, even if we avoid using threads ourselves: OSGi is implicitly multi-threaded. Therefore we must ensure that our code is appropriately thread-safe, particularly when receiving events or callbacks from the framework or from other bundles.

This chapter is an introduction to the concepts of concurrency in Java as applied to OSGi. For a more thorough treatment of Java concurrency, “Java Concurrency in Practice” by Brian Goetz et al [15] is invaluable and, in the author's opinion, should be kept close at hand by all professional Java programmers.

6.1 The Price of Freedom

Let's follow through a simple scenario involving the imaginary bundles *A*, *B* and *C*, which is illustrated in the style of a UML sequence diagram in Figure 6.1. Suppose bundle *A* starts a thread and at some point it obtains a handle to bundle *B* and calls the `start` method from that thread. This will cause bundle *B* to activate, and the `start` method of *B*'s activator will be invoked in the thread created by *A*. Furthermore, suppose that *B* registers a service during its `start` method and *C* happens to be listening with a service tracker for instances of that service type. The `addingService` method of *C*'s tracker will be called, also from the thread created by *A*. Finally, suppose *C* creates

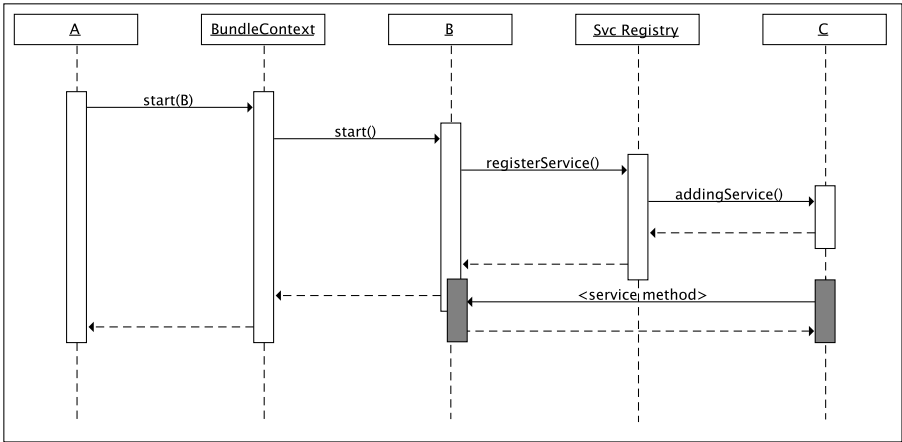


Figure 6.1: Framework Calls and Callbacks in OSGi

a polling thread which periodically calls the service registered by *B*. The methods of *B*'s service will be executed in a thread created by *C*.

When a client gets a service from the service registry, it sees the real service object, not any kind of proxy or wrapper. Therefore when the client invokes methods on the service, those invocations are standard synchronous method calls, which means that the service method executes on a thread which is “owned” by the client bundle.

Furthermore, many notifications in OSGi (but not all) happen *synchronously*. That is, when the framework is called with a method that generates callbacks — such as sending a `ServiceEvent` to registered service listeners, or calling a bundle activator’s `start` or `stop` methods — those callbacks are executed in the same thread, and must complete before control is returned to the caller of the framework method.

The above has three major implications:

- Callbacks and service methods can and will be invoked from *any arbitrary thread*, perhaps even from many threads at the same time. Our code will fail in unpredictable ways unless we are conscious of this and allow for it.
- The thread that our callbacks and service methods are invoked from does not “belong” to us. If we execute long-running operations or blocking I/O calls directly from these callbacks then we can delay the entire system.
- When we call an OSGi API method we can’t predict the set of callbacks and listeners that will be called by the framework as a result, and therefore what locks those callbacks will attempt to take. If we hold other

locks while calling such methods, we risk causing a deadlock.

The only solution to these problems is to use good concurrent programming practices. However, safe concurrency is really not that difficult, at least in the sense of being intellectually challenging, like quantum physics or a game of Go. One does not need to be a genius to do it correctly. The key is *discipline* — something that many geniuses lack! As long as we consistently apply a few simple rules, we can easily handle most situations we encounter.

1. Immutable objects are automatically thread-safe, and objects that are not shared across threads do not need to be thread-safe. Therefore share as little as possible and favour immutability wherever possible.
2. When objects really must be shared and mutable, guard all accesses (both read *and* write) to shared fields with a lock on the same object, or make appropriate use of volatile variables.
3. Avoid acquiring new locks when holding an existing lock. As a direct consequence of this rule, we must avoid holding any locks when calling unknown or “foreign” code that might attempt to acquire a lock. This includes calls to services or to OSGi APIs, many of which can result in callbacks to other bundles that execute in *our* thread.

6.2 Shared Mutable State

We can substantially reduce the size of the concurrency problem simply by sharing as few objects as possible, and making as many shared objects as possible immutable. Unfortunately it’s generally not possible to reduce the problem size to zero by these methods. In most real-world applications we cannot completely avoid sharing mutable state, so we need to find a way to do it safely.

Let’s look at some example code based on services. Suppose we wish to maintain a `Map` of registered mailbox services, keyed by name, and we wish to offer a public method `getMailboxByName` which will return the specified mailbox, or `null` if none such currently exists. The sample code in Listing 6.1 uses traditional Java synchronized blocks to achieve thread safety. This class is well behaved because it follows the rules: all accesses to the map, *including* simple read operations, are made from the context of a lock on the same object. Although the map field itself is final, the *content* of the map is mutable, and it is shared across threads, so it needs to be protected by a lock. Also, the synchronized blocks are as short as possible, quickly releasing the lock when it is no longer required.

However, the code in Listing 6.1 does not take any advantage of the new concurrency features introduced in Java 5. Using those features, we can do slightly

Listing 6.1 Thread Safety using Synchronized Blocks

```
1 package org.osgi.book.reader.mailboxmap;
2
3 import java.util.HashMap;
4 import java.util.Map;
5
6 import org.osgi.book.reader.api.Mailbox;
7 import org.osgi.framework.BundleActivator;
8 import org.osgi.framework.BundleContext;
9 import org.osgi.framework.ServiceReference;
10 import org.osgi.util.tracker.ServiceTracker;
11
12 public class MailboxMapActivator1 implements BundleActivator {
13
14     private final Map map = new HashMap();
15     private ServiceTracker tracker;
16
17     public void start(BundleContext context) throws Exception {
18         tracker = new MapTracker(context);
19         tracker.open();
20     }
21     public void stop(BundleContext context) throws Exception {
22         tracker.close();
23     }
24
25     public Mailbox getMailboxByName(String name) {
26         synchronized (map) {
27             return (Mailbox) map.get(name);
28         }
29     }
30
31     private class MapTracker extends ServiceTracker {
32
33         public MapTracker(BundleContext context) {
34             super(context, Mailbox.class.getName(), null);
35         }
36
37         public Object addingService(ServiceReference reference) {
38             String mboxName = (String) reference
39                 .getProperty(Mailbox.NAME_PROPERTY);
40             Mailbox mbox = (Mailbox) context.getService(reference);
41             synchronized (map) {
42                 map.put(mboxName, mbox);
43             }
44             return mboxName;
45         }
46
47         public void removedService(ServiceReference reference,
48             Object service) {
49             String mboxName = (String) service;
50             synchronized (map) {
51                 map.remove(mboxName);
52             }
53         }
54     }
55 }
```

better thanks to the observation that many threads should be able to call the `getMailboxByName` method at the same time. Traditional synchronized blocks cannot distinguish between read operations and write operations: they ensure *all* operations inside them have exclusive access to the lock object. However Java 5 introduced Read/Write locks that do make this distinction. Listing 6.2 uses this feature to allow many threads to read from the map concurrently, while still preventing concurrent updates to the map:

If Java 5 (or above) is available to you, then it is well worth investigating whether the new concurrency library therein will be of benefit. However even if you are limited to Java 1.4 or earlier, you should not fear the synchronized block. Many developers avoid synchronization because of problems related to its performance, but it is really not that bad, especially when locks are un-contended — this can usually be achieved by sticking to fine-grained synchronized blocks. There is no sense in sacrificing correct concurrent behaviour for a minor performance gain.

6.3 Safe Publication

Safe publication means making an object available to be accessed by other threads so that those threads do not see the object in an invalid or partially constructed state.

“Publishing” an object entails placing it in a location from which it can be read by another thread. The simplest example of such a location is a public field of an already-published object, but placing the object in a private field also effectively publishes it if there is a public method accessing that field. Unfortunately this mechanism is not enough on its own for the publication to be *safe*.

Listing 6.3 is quoted from [15] as an example of unsafe publication. Because of the way modern CPU architectures work, and the way those architectures are exposed to Java programs through the Java Memory Model, a thread reading the `holder` field might see a `null` value or it might see a partially constructed version of the `Holder` object, even if that thread reads the variable *after* `initialize` was called in another thread.

There are four ways to safely publish an object, as listed in [15]:

- Initialise an object reference from a static initialiser.
- Store a reference into a `volatile` field or `AtomicReference`.
- Store a reference into a `final` field of a properly constructed object.
- Store a reference into a field that is properly guarded by a lock, i.e., making appropriate use of `synchronized` blocks or methods.

Listing 6.2 Thread Safety using Read/Write Locks

```
1 package org.osgi.book.reader.mailboxmap;
2
3 import java.util.HashMap;
4 import java.util.Map;
5 import java.util.concurrent.locks.ReadWriteLock;
6 import java.util.concurrent.locks.ReentrantReadWriteLock;
7
8 import org.osgi.book.reader.api.Mailbox;
9 import org.osgi.framework.BundleActivator;
10 import org.osgi.framework.BundleContext;
11 import org.osgi.framework.ServiceReference;
12 import org.osgi.util.tracker.ServiceTracker;
13
14 public class MailboxMapActivator2 implements BundleActivator {
15
16     private final Map<String, Mailbox> map =
17         new HashMap<String, Mailbox>();
18     private final ReadWriteLock mapLock = new ReentrantReadWriteLock();
19     private volatile ServiceTracker tracker;
20
21     public void start(BundleContext context) throws Exception {
22         tracker = new MapTracker(context);
23         tracker.open();
24     }
25     public void stop(BundleContext context) throws Exception {
26         tracker.close();
27     }
28     public Mailbox getMailboxByName(String name) {
29         try {
30             mapLock.readLock().lock();
31             return map.get(name);
32         } finally {
33             mapLock.readLock().unlock();
34         }
35     }
36
37     private class MapTracker extends ServiceTracker {
38
39         public MapTracker(BundleContext context) {
40             super(context, Mailbox.class.getName(), null);
41         }
42         public Object addingService(ServiceReference reference) {
43             String mboxName = (String) reference
44                 .getProperty(Mailbox.NAME_PROPERTY);
45             Mailbox mbox = (Mailbox) context.getService(reference);
46             try {
47                 mapLock.writeLock().lock();
48                 map.put(mboxName, mbox);
49             } finally {
50                 mapLock.writeLock().unlock();
51             }
52             return mboxName;
53         }
54         public void removedService(ServiceReference reference,
55             Object service) {
56             String mboxName = (String) service;
57             try {
58                 mapLock.writeLock().lock();
59                 map.remove(mboxName);
60             } finally {
61                 mapLock.writeLock().unlock();
62             }
63         }
64     }
65 }
```

Listing 6.3 Unsafe Publication

```
// Unsafe publication
public Holder holder;

public void initialize() {
    holder = new Holder(42);
}
```

6.3.1 Safe Publication in Services

Suppose we have a “dictionary” service that allows us to both look up definitions of words and add new definitions. The very simple interface for this service is shown in Listing 6.4 and a trivial (but broken) implementation is in Listing 6.5.

Listing 6.4 Dictionary Service interface

```
1 package org.osgi.book.concurrency;
2
3 public interface DictionaryService {
4     void addDefinition(String word, String definition);
5     String lookup(String word);
6 }
```

Listing 6.5 Unsafe Publication in a Service

```
1 package org.osgi.book.concurrency;
2
3 import java.util.HashMap;
4 import java.util.Map;
5
6 public class UnsafeDictionaryService implements DictionaryService {
7
8     private Map<String,String> map;
9
10    public void addDefinition(String word, String definition) {
11        if(map == null) map = new HashMap<String, String>();
12
13        map.put(word, definition);
14    }
15
16    public String lookup(String word) {
17        return map == null ? null : map.get(word);
18    }
19 }
```

Recall that a service, once registered, can be called from any thread at any time. The `map` field is only initialised on first use of the `addDefinition` method, but a call to the `lookup` method could see a partially constructed `HashMap` object, with unpredictable results¹.

¹There are other concurrency problems here too. For example, two threads could enter

The root of the problem here is clearly the late initialisation of the `HashMap` object, which appears to be a classic example of “premature optimisation”. It would be easier and safer to create the map during construction of our service object. We could then mark it `final` to ensure it is safely published. This is shown in Listing 6.6. Note that we also have to use a thread-safe `Map` implementation rather than plain `HashMap` — we can get one by calling `Collections.synchronizedMap` but we could also have used a `ConcurrentHashMap`.

Listing 6.6 Safe Publication in a Service

```

1 package org.osgi.book.concurrency;
2
3 import java.util.Collections;
4 import java.util.HashMap;
5 import java.util.Map;
6
7 public class SafeDictionaryService implements DictionaryService {
8
9     private final Map<String,String> map =
10         Collections.synchronizedMap(new HashMap<String,String>());
11
12     public void addDefinition(String word, String definition) {
13         map.put(word, definition);
14     }
15
16     public String lookup(String word) {
17         return map.get(word);
18     }
19 }

```

Using a `final` field is the preferred way to safely publish an object, since it results in code that is easy to reason about, but unfortunately this is not always possible. We may not be able to create the object we want during the construction of our service object if it has a dependency on an object passed by a client of the service. For example, take a look at the service interface in Listing 6.7 and the unsafe implementation in Listing 6.8. The problem is now clear but we cannot fix it by moving initialisation of the `connection` field to the constructor and making it `final`, since it depends on the `DataSource` object passed by a client².

A solution to this problem is to simply declare the `connection` field to be `volatile`. This means that any modifications will be automatically visible in

`addDefinition` at about the same time and both see a `null` value for the `map` field, and so they would both create a new `HashMap` and put the word definition into it. But only one `HashMap` instance would ultimately remain, so one of the definitions would be lost.

²Note that there are more problems with this service. The `initialise` method is unsafe because two clients could call it at the same time. In fact the service interface itself is poorly designed: suppose a client calls `initialise` and then the service instance it called goes away and is replaced by an alternative — the client would need to track the change and call `initialise` again. For this reason we should avoid designing services that require conversational state, i.e., a controlled series of method calls and responses. Services should ideally be stateless.

Listing 6.7 Connection Cache interface

```
1 package org.osgi.book.concurrency;
2
3 import java.sql.Connection;
4 import java.sql.SQLException;
5
6 import javax.sql.DataSource;
7
8 public interface ConnectionCache {
9     void initialise(DataSource dataSource) throws SQLException;
10    Connection getConnection();
11 }
```

Listing 6.8 Unsafe Connection Cache

```
1 package org.osgi.book.concurrency;
2
3 import java.sql.Connection;
4 import java.sql.SQLException;
5
6 import javax.sql.DataSource;
7
8 public class UnsafeConnectionCache implements ConnectionCache {
9
10    private Connection connection;
11
12    public void initialise(DataSource dataSource) throws SQLException {
13        connection = dataSource.getConnection();
14    }
15
16    public Connection getConnection() {
17        return connection;
18    }
19 }
```

full to any other threads, so we can safely publish an object simply by assigning it into a `volatile` field. Note that this behaviour is only guaranteed on Java 5 and above — developers working with older Java versions will need to use a `synchronized` block or method.

6.3.2 Safe Publication in Framework Callbacks

Given the warnings above, we might expect that the activator class in Listing 6.9 would not be thread-safe. It appears to have a serious problem with unsafe publication of the `LogTracker` object into the `log` field, which is declared neither `final` nor `volatile`.

Listing 6.9 Is This Bundle Activator Thread-safe?

```

1 package org.osgi.book.concurrency;

3 import org.osgi.book.utils.LogTracker;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6 import org.osgi.framework.BundleEvent;
7 import org.osgi.framework.SynchronousBundleListener;
8 import org.osgi.service.log.LogService;

10 public class DubiousBundleActivator implements BundleActivator,
11     SynchronousBundleListener {

13     private LogTracker log;

15     public void start(BundleContext context) throws Exception {
16         log = new LogTracker(context);
17         log.open();

19         context.addBundleListener(this);
20     }

22     public void bundleChanged(BundleEvent event) {
23         if(BundleEvent.INSTALLED == event.getType()) {
24             log.log(LogService.LOG_INFO, "Bundle installed");
25         } else if(BundleEvent.UNINSTALLED == event.getType()) {
26             log.log(LogService.LOG_INFO, "Bundle removed");
27         }
28     }

30     public void stop(BundleContext context) throws Exception {
31         context.removeBundleListener(this);
32         log.close();
33     }
34 }

```

If the problem is not yet apparent, consider that all three methods of this class are technically *callbacks*. The `start` callback happens when a bundle decides to explicitly request the framework to start our bundle, and it runs in the calling thread of that bundle. The `stop` callback likewise happens when a bundle (not necessarily the same one) decides to request the framework to stop our bundle, and it also runs in the calling thread of that bundle. The

`bundleChanged` callback is called whenever any bundle changes its state while we have our activator registered as a listener for bundle events. Therefore all three methods can be called on different threads, so we should use safe publication for the `LogTracker` object accessed from each one.

In fact, it turns out that Listing 6.9 is perfectly thread-safe! But the reasons why it is thread-safe are subtle and require some understanding of how the OSGi framework itself is implemented, and for that reason it is this author's opinion that developers should *use safe publication idioms anyway*.

These two statements clearly need some explanation! First, the reason why Listing 6.9 is thread-safe. One of the ways that we can achieve safe publication, besides `final` and `volatile`, is to use a `synchronized` block or method. In the terminology of the Java Memory Model, we need to ensure that the publication of the `LogTracker` object into the `log` field “*happens-before*” the read operations on that field, and the Java Memory model also guarantees that everything that happens on a thread prior to releasing a lock (i.e., exiting from a `synchronized` block) “*happens-before*” anything that happens on another thread after that thread acquires the *same* lock. So by synchronising on an object before setting the `log` field and before accessing it, we can achieve safe publication of the `LogTracker` object it contains. But crucially, it does not matter *which* object is used to synchronise, so long as the same one is consistently used.

Now, the OSGi framework uses a lot of synchronisation internally — it must, because it is able to offer most of its features safely to multiple threads — and we can use this knowledge to “piggyback” on existing synchronisation in the framework, meaning that in many cases we don't need to add our own safe publication idioms such as `final`, `volatile` or `synchronized`. The trick is knowing when we can get away with this and when we cannot.

So for example the framework holds, somewhere in its internal memory, a list of bundle listeners. When we call `BundleContext.addBundleListener` on line 19, the framework uses a `synchronized` block to add us to the list. Later when a bundle state changes, the framework will use another `synchronized` block to get a snapshot of the currently installed listeners before calling `bundleChanged` on those listeners. Therefore everything we did before calling `addBundleListener` *happens-before* everything we do after being called in `bundleChanged`, and our access of the `log` field in lines 24 and 26 is safe.

Similarly, the framework holds the state of all bundles. When our `start` method completes, the framework changes the state of the bundle from `STARTING` to `ACTIVE`, and it uses a `synchronized` block to do that. It is illegal to call `stop` on a bundle that is not `ACTIVE` (or rather, such a call will be ignored), so the framework will check the state of our bundle (using a `synchronized` block) before allowing our `stop` method to run. Therefore another *happens-before* relationship exists which we can exploit, making our access of

the `log` field safe on line 32 also.

We can express these deduced relationships (plus another that happens to exist) more succinctly as “rules” similar to the built-in rules offered by the Java Memory Model itself:

Bundle Activator rule. Each action in the `start` method of a `BundleActivator` *happens-before* every action in the `stop` method.

Listener registration rule. Registering a listener with the framework — including `ServiceListener`, `BundleListener`, `SynchronousBundleListener` or `FrameworkListener` — *happens-before* any callback is invoked on that listener.

Service registration rule. Registering a service *happens-before* any invocation of the methods of that service by a client.

Unfortunately there are problems with relying on these “rules”. First, they are not true rules written down in the OSGi specification such that all conforming implementations must abide by them³. They are merely deduced from “circumstantial evidence”, i.e., the multi-threadedness of the OSGi framework and our expectation of how that is handled internally by the framework. This raises the possibility that a conforming OSGi implementation may perform locking in a slightly different way, so violating these rules.

But the bigger problem is that piggybacking on the framework’s synchronisation is a very advanced technique. It just happens that we can write “dumb” code such as the code in Listing 6.9 and find that it works... but we must be very aware of when our license runs out and we must start caring about safe publication again. To do this, we need an intimate understanding of internal framework issues. It seems it would be easier just to use safe publication idioms everywhere. In this example, that simply means adding the `volatile` keyword to the `log` field.

6.4 Don’t Hold Locks when Calling Foreign Code

Suppose we have a service that implements the `MailboxRegistrationService` interface shown in Listing 6.10.

When the `registerMailbox` method is called, we should register the supplied mailbox as a service with the specified name. Also we should unregister any previous mailbox service that we may have registered with that same name.

Listing 6.11 shows a naïve implementation of this service interface. This code has a problem: the whole `registerMailbox` method is synchronized, including

³Though perhaps they should be!

Listing 6.10 Mailbox Registration Service Interface

```

1 package org.osgi.book.concurrency;
2
3 import org.osgi.book.reader.api.Mailbox;
4
5 public interface MailboxRegistrationService {
6     void registerMailbox(String name, Mailbox mailbox);
7 }

```

the calls to OSGi to register and unregister the services. That means we will enter the OSGi API with a lock held and therefore any callbacks — such as the `addingService` method of a tracker — will also be called with that lock held. If any of those callbacks then try to acquire another lock, we may end up in deadlock.

Listing 6.11 Holding a lock while calling OSGi APIs

```

11 public class BadLockingMailboxRegistrationService implements
12     MailboxRegistrationService {
13
14     private final Map<String, ServiceRegistration> map
15         = new HashMap<String, ServiceRegistration>();
16     private final BundleContext context;
17
18     public BadLockingMailboxRegistrationService(BundleContext context) {
19         this.context = context;
20     }
21
22     // DO NOT DO THIS!
23     public synchronized void registerMailbox(String name,
24         Mailbox mailbox) {
25         ServiceRegistration priorReg = map.get(name);
26         priorReg.unregister();
27
28         Properties props = new Properties();
29         props.put(Mailbox.NAME_PROPERTY, name);
30         ServiceRegistration reg = context.registerService(
31             Mailbox.class.getName(), mailbox, props);
32
33         map.put(name, reg);
34     }
35 }

```

The Wikipedia definition of *Deadlock*^[16] refers to a charmingly illogical extract from an Act of the Kansas State Legislature:

“When two trains approach each other at a crossing, both shall come to a full stop and neither shall start up again until the other has gone.”^[17]

This is a classic deadlock: neither train can make any progress because it is waiting for the other train to act first.

Another familiar example of deadlock is the so-called “dining philosophers”




Figure 6.2: The Dining Philosophers Problem, Simplified

problem, first described by Edsger Dijkstra in [18]. To reduce this problem to its bare minimum, imagine two philosophers at a dinner table waiting to eat, but there is only one knife and one fork (Figure 6.2). A philosopher needs both utensils in order to eat, but one immediately picks up the knife the other immediately picks up the fork. They both then wait for the other utensil to become available. It should be clear that this strategy will eventually result in both philosophers starving to death.

To translate this unhappy situation to OSGi programming, imagine that a thread has taken a lock on an object F . Then it tries to call our `registerMailbox`, which locks object K — but it must wait, perhaps because another thread is already executing `registerMailbox`. One of the callbacks resulting from the service registration then attempts to lock F . The result: two starving threads and no work done.

The traditional solution to the dining philosophers problem is simply to always take the knife and fork in the same order. If both philosophers try to get the fork before the knife, then the first one to pick up the fork will eat first, and then the second philosopher will eat when the first is finished. This may not be very fair, but at least nobody starves. Similarly, the safe way to take locks on multiple objects is to always take them in the same order. The ordering can be arbitrary but the important thing is to consistently apply whatever ordering scheme is chosen (the philosophers also get to eat if they both attempt to take the knife first — it is only when they choose a different first utensil that they starve).

Unfortunately when calling an OSGi API method that involves callbacks into

other bundles, we simply cannot enforce any such ordering, because those bundles cannot know about our ordering scheme. The only alternative is to restructure our code to avoid holding any locks when calling the OSGi APIs. Note that avoiding locks does not just mean avoiding the `synchronized` keyword. We could achieve the level of isolation we desire in the `registerMailbox` method by using a binary semaphore or an instance of the `Lock` classes. However when used in such a way, these constructs have exactly the same semantics as a `synchronized` method or block, and can result in the same types of deadlock.

The trick is avoid holding locks during calls to OSGi APIs, but to do this *without* losing the atomicity guarantees that we require. Sometimes this requires some re-ordering of the operations. For example, Listing 6.12 shows a version of the service that uses a lock only to manipulate the `map` field⁴. The result of the operation passes out of the locked region and tells us whether there was a prior registration of the service which needs to be unregistered. Note that the `put` method of `Map` returns the previous mapping for the key if one existed.

Due to the reordering there is a slight change in the externally observable behaviour of this method. The previous version (if it worked!) would result in a short gap between the old service being removed and the new service being created, during which there would be no service available. The new version reverses the steps: the new service is created very slightly *before* the old service is removed, so there will briefly be two services rather than none. It turns out this is not such a bad thing: as well as making it possible to reduce the size of our locked regions, consumers of the service also benefit from having a replacement service immediately available when the first service goes away.

6.5 GUI Development

Another area where multi-threading causes pain is in programming graphical user interfaces (GUIs). Almost all GUI libraries — including the most popular Java ones, Swing and SWT — insist that all calls to those libraries must be made from a single thread, the “event dispatch thread” or EDT. But as we have seen, when our callbacks or service methods are executed, we have no idea whether we are in the EDT or some other, arbitrary thread. Therefore we have to use utilities supplied by the GUI library to pass blocks of code that will be executed in the EDT when it gets around to it. In Swing, we need to pass a `Runnable` instance to the `SwingUtilities.invokeLater` method, but for efficiency, we should first check whether we’re already in the EDT by calling `EventQueue.isDispatchThread`

⁴In this example we use a plain `HashMap` and wrap the calls to it in a `synchronized` block in order to be explicit about the locking. We could have used a `ConcurrentHashMap` which performs fine-grained internal locking with no need for a `synchronized` block.

Listing 6.12 Avoiding holding a lock while calling OSGi APIs

```
1 package org.osgi.book.concurrency;
2
3 import java.util.HashMap;
4 import java.util.Map;
5 import java.util.Properties;
6
7 import org.osgi.book.reader.api.Mailbox;
8 import org.osgi.framework.BundleContext;
9 import org.osgi.framework.ServiceRegistration;
10
11 public class GoodLockingMailboxRegistrationService implements
12     MailboxRegistrationService {
13
14     private final Map<String, ServiceRegistration> map
15         = new HashMap<String, ServiceRegistration>();
16
17     private final BundleContext context;
18
19     public GoodLockingMailboxRegistrationService(BundleContext context) {
20         this.context = context;
21     }
22
23     public void registerMailbox(String name, Mailbox mailbox) {
24         Properties props = new Properties();
25         props.put(Mailbox.NAME_PROPERTY, name);
26         ServiceRegistration reg = context.registerService(
27             Mailbox.class.getName(), mailbox, props);
28
29         ServiceRegistration priorReg;
30         synchronized (map) {
31             priorReg = map.put(name, reg);
32         }
33
34         if (priorReg != null) {
35             priorReg.unregister();
36         }
37     }
38 }
```

The example in Listing 6.13 uses this technique to update the text of a Swing label object to indicate the current number of mailbox services.

Unfortunately, programming with Swing (or any GUI library) in a multi-threaded environment can quickly become cumbersome due to the need to create anonymous inner classes implementing the `Runnable` interface. We cannot refactor the complexity out into library methods because the Java language does not support closures or “blocks”. This is an area where we can benefit from using a more powerful programming language, such as Scala or JRuby. For example, assuming we have defined a Scala library function for performing actions in the EDT as shown in Listing 6.14, we can then write a same tracker class in Scala using the code in Listing 6.15.

6.6 Using Executors

Perversely, despite the cumbersome code, working with GUIs can *simplify* some aspects of concurrent programming. Since we always have to transfer work to be run on a single thread, we can take advantage of that thread to perform mutations to state without needing to take any locks at all — so long as the variables we mutate are *only* touched within the GUI thread.

We can exploit this pattern even when not writing GUIs. In Section 6.4 we saw the need to avoid holding locks when making calls to certain OSGi API methods, and in the example code we had to slightly reorder the operations to allow us to safely perform those OSGi calls outside of a lock. This led to a small change in the behaviour of the method — brief periods of time in which two Mailbox services are registered.

Suppose we are under some constraint which means we simply cannot have two Mailbox services registered at the same time. This means we *must* first unregister the existing service before registering the new one. Now we have a problem: without wrapping the two operations in a synchronized block, we cannot make them atomic, so a race condition could occur in which two threads both simultaneously unregister the existing service and then both create and register new services... the very situation we were trying to avoid!

There is a solution which satisfies both the requirement not to lock and the requirement to update atomically: simply perform all of the updates from a single thread. We could do this by handing updates to a thread that we create specifically for the purpose of updating our Mailbox service. In Java 5 this is easily done with an `Executor` as shown in Listing 6.16.

Of course, this solution also produces a subtle change to the behaviour of the service. Now the update to the service doesn't happen synchronously during the call to `registerMailbox`, but asynchronously shortly afterwards.

Listing 6.13 Updating a Swing Control in Response to a Service Event

```
1 package org.osgi.book.reader.tutorial;
2
3 import java.awt.EventQueue;
4
5 import javax.swing.JLabel;
6 import javax.swing.SwingUtilities;
7
8 import org.osgi.book.reader.api.Mailbox;
9 import org.osgi.framework.BundleContext;
10 import org.osgi.framework.ServiceReference;
11 import org.osgi.util.tracker.ServiceTracker;
12
13 public class JLabelMailboxCountTracker extends ServiceTracker {
14
15     private final JLabel label;
16     private int count = 0;
17
18     public JLabelMailboxCountTracker(JLabel label,
19         BundleContext context) {
20         super(context, Mailbox.class.getName(), null);
21         this.label = label;
22     }
23
24     @Override
25     public Object addingService(ServiceReference reference) {
26         int displayCount;
27         synchronized (this) {
28             count++;
29             displayCount = count;
30         }
31         updateDisplay(displayCount);
32         return null;
33     }
34
35     @Override
36     public void removedService(ServiceReference reference,
37         Object service) {
38         int displayCount;
39         synchronized (this) {
40             count--;
41             displayCount = count;
42         }
43         updateDisplay(displayCount);
44     }
45
46     private void updateDisplay(final int displayCount) {
47         Runnable action = new Runnable() {
48             public void run() {
49                 label.setText("There are " + displayCount + " mailboxes");
50             }
51         };
52         if(EventQueue.isDispatchThread()) {
53             action.run();
54         } else {
55             SwingUtilities.invokeLater(action);
56         }
57     }
58 }
```

Listing 6.14 A Scala Utility to Execute a Closure in the Event Thread

```
def inEDT(action: => Unit) =
  if (EventQueue.isDispatchThread())
    action
  else
    SwingUtilities.invokeLater(new Runnable() {def run = action})
```

Listing 6.15 Updating a Swing Control — Scala Version

```
class JLabelMailboxCountTracker(ctx: BundleContext, label: JLabel)
  extends ServiceTracker(ctx, classOf[Mailbox].getName, null) {

  private var count = 0: int

  override
  def addingService(ref: ServiceReference): AnyRef = {
    var displayCount = 0
    synchronized { count += 1; displayCount = count }

    inEDT(label.setText("foo"))
    null
  }

  override
  def removedService(ref: ServiceReference, service: AnyRef) = {
    var displayCount = 0
    synchronized { count -= 1; displayCount = count }

    inEDT(label.setText("foo"))
  }
}
```

Listing 6.16 Single-threaded execution

```
1 package org.osgi.book.concurrency;

3 import java.util.HashMap;
4 import java.util.Map;
5 import java.util.Properties;
6 import java.util.concurrent.ExecutorService;
7 import java.util.concurrent.Executors;

9 import org.osgi.book.reader.api.Mailbox;
10 import org.osgi.framework.BundleContext;
11 import org.osgi.framework.ServiceRegistration;

13 public class SingleThreadedMailboxRegistrationService implements
14     MailboxRegistrationService {

16     private final Map<String, ServiceRegistration> map
17         = new HashMap<String, ServiceRegistration>();
18     private final BundleContext context;

20     private final ExecutorService executor =
21         Executors.newSingleThreadExecutor();

23     public SingleThreadedMailboxRegistrationService(
24         BundleContext context) {
25         this.context = context;
26     }

28     public void registerMailbox(final String name,
29                               final Mailbox mailbox) {
30         Runnable task = new Runnable() {
31             public void run() {
32                 ServiceRegistration priorReg = map.get(name);
33                 priorReg.unregister();

35                 Properties props = new Properties();
36                 props.put(Mailbox.NAME_PROPERTY, name);
37                 ServiceRegistration reg = context.registerService(
38                     Mailbox.class.getName(), mailbox, props);
39                 map.put(name, reg);
40             }
41         };
42         executor.execute(task);
43     }

45     public void cleanup() {
46         executor.shutdown();
47     }
48 }
```

This is essentially a trade-off that we cannot escape — the requirement to avoid holding a lock forces us to either reorder operations (as in the previous solution) or execute asynchronously. In most cases the reordering solution is preferable.

There is another problem with the code in Listing 6.16: it always creates its own thread irrespective of any application-wide management policies with respect to thread creation. Threads are somewhat expensive in Java, and if too many services create their own single-purpose threads then we will have a problem managing our application. Note also that we need to remember to cleanup this service to ensure the thread is shutdown when no longer needed.

In this case we can employ a useful form of executor called the `SerialExecutor`, the code for which appears in Listing 6.17. This executor ensures that at most one task is running at any time, but it achieves this without creating a thread or locking anything during the execution of a task. The trick is an internal *work queue*. If a call to `execute` arrives while the executor is idle, then the thread making that call becomes the “master” and it immediately executes the task... note that the task is executed *synchronously* in the calling thread. However if a call to `execute` arrives while a master is already running then we simply add the task to the end of the work queue and immediately return to the caller. The task will be executed by the master thread, which ensures that the work queue is emptied before returning to its caller. As soon as the master thread returns from `execute`, the executor is idle and the next thread to call `execute` will become the new master.

Listing 6.18 shows the Mailbox registration service yet again, using `SerialExecutor`. In this version we can get rid of the `cleanup` method since there is no thread to shutdown, but we need to switch the map to a thread-safe version for visibility, since any thread can become the master thread.

The `SerialExecutor` class is useful, and you should consider keeping it handy in your OSGi toolbox, but it is also not a panacea. We must still consider it asynchronous, since non-master threads will return to their callers before the work is completed by the master thread. Also it is completely unsuitable when calls to `execute` are made with too high frequency, since any thread unlucky enough to be nominated the master will be stuck executing the work of many other threads before it is finally able to return to its caller. A possible enhancement would be to allow `SerialExecutor` to spin off a thread if the master thread decided that enough was enough.

Another useful pattern is to have a bundle export one or more executors as services. In doing this we resolve the dilemma of having bundles that wish to spin off threads to perform work asynchronously versus the desire to manage creation of threads at an application level. We can create a single “work manager” bundle that creates thread pools in various configurations — single threads, fixed size pools, resizable pools, etc. — and then any bundle wishing

Listing 6.17 The `SerialExecutor` class

```

1 package org.osgi.book.utils;

3 import java.util.ArrayList;
4 import java.util.List;
5 import java.util.concurrent.Executor;

7 import org.osgi.service.log.LogService;

9 public class SerialExecutor implements Executor {
10     private final List<Runnable> queue = new ArrayList<Runnable>();
11     private final ThreadLocal<Integer> taskCount;
12     private final LogService log;

14     public SerialExecutor() {
15         this(null);
16     }

18     public SerialExecutor(LogService log) {
19         this.log = log;
20         taskCount = new ThreadLocal<Integer>() {
21             protected Integer initialValue() {
22                 return 0;
23             }
24         };
25     }

27     public void execute(Runnable command) {
28         Runnable next = null;
29         int worked = 0;

31         // If the queue is empty, I am the master and my next task is to
32         // execute my own work. If the queue is non-empty, then I simply
33         // add my task to the queue and return.
34         synchronized (this) {
35             next = queue.isEmpty() ? command : null;
36             queue.add(command);
37         }

39         while (next != null) {
40             // Do the work!
41             try {
42                 next.run();
43                 worked++;
44             } catch (Exception e) {
45                 logError("Error processing task", e);
46             }

48             // Get more work if it exists on the queue
49             synchronized (this) {
50                 queue.remove(0); // Head element is the one just processed
51                 next = queue.isEmpty() ? null : queue.get(0);
52             }

54             taskCount.set(worked);
55         }
56     }

```

Listing 6.17 (continued)

```

58  /**
59   * Returns the number of tasks executed by the last call to
60   * <code>execute</code> from the calling thread.
61   */
62  public int getLastTaskCount() {
63      return taskCount.get();
64  }

66  private void logError(String message, Exception e) {
67      if (log != null) {
68          log.log(LogService.LOG_ERROR, message, e);
69      }
70  }
71 }

```

Listing 6.18 The Mailbox registration service using SerialExecutor

```

1  package org.osgi.book.concurrency;

3  import java.util.HashMap;
4  import java.util.Map;
5  import java.util.Properties;
6  import java.util.concurrent.ConcurrentHashMap;
7  import java.util.concurrent.Executor;

9  import org.osgi.book.reader.api.Mailbox;
10 import org.osgi.book.utils.SerialExecutor;
11 import org.osgi.framework.BundleContext;
12 import org.osgi.framework.ServiceRegistration;

14 public class SerialMailboxRegistrationService implements
15     MailboxRegistrationService {

17     private final Map<String, ServiceRegistration> map =
18         new ConcurrentHashMap<String, ServiceRegistration>();
19     private final BundleContext context;

21     private final Executor executor = new SerialExecutor();

23     public SerialMailboxRegistrationService(BundleContext context) {
24         this.context = context;
25     }

27     public void registerMailbox(final String name,
28         final Mailbox mailbox) {
29         Runnable task = new Runnable() {
30             public void run() {
31                 ServiceRegistration priorReg = map.get(name);
32                 priorReg.unregister();

34                 Properties props = new Properties();
35                 props.put(Mailbox.NAME_PROPERTY, name);
36                 ServiceRegistration reg = context.registerService(
37                     Mailbox.class.getName(), mailbox, props);

39                 map.put(name, reg);
40             }
41         };
42         executor.execute(task);
43     }
44 }

```

to execute tasks can submit them via the service interface. Listing 6.19 shows an example of registering such an executor; note that we register a wrapper object around the thread pool rather than registering the thread pool directly, as this prevents clients from casting the service to its `ExecutorService` interface and calling methods such as `shutdown` on it.

Listing 6.19 Registering a Thread Pool as an Executor Service

```

1 package org.osgi.book.concurrency;

3 import java.util.Properties;
4 import java.util.concurrent.Executor;
5 import java.util.concurrent.ExecutorService;
6 import java.util.concurrent.Executors;

8 import org.osgi.framework.BundleActivator;
9 import org.osgi.framework.BundleContext;
10 import org.osgi.framework.ServiceRegistration;

12 public class ExecutorServiceActivator implements BundleActivator {

14     private static final int POOL_SIZE = 20;

16     private volatile ExecutorService threadPool;
17     private volatile ServiceRegistration svcReg;

19     public void start(BundleContext context) {
20         threadPool = Executors.newFixedThreadPool(POOL_SIZE);

22         Executor wrapper = new Executor() {
23             public void execute(Runnable command) {
24                 threadPool.execute(command);
25             }
26         };

28         Properties props = new Properties();
29         props.put("kind", "fixed");
30         props.put("poolSize", POOL_SIZE);
31         svcReg = context.registerService(Executor.class.getName(),
32             wrapper, props);
33     }

35     public void stop(BundleContext context) {
36         svcReg.unregister();
37         threadPool.shutdown();
38     }
39 }

```

6.7 Interrupting Threads

As we discussed in Chapter 2, there is no general way to force a thread to stop in Java, except by shutting down the whole Java Virtual Machine. This is by design: if such a mechanism existed then threads could be terminated while mutating data structures, and would be likely to leave them in an inconsistent state. Instead, Java offers a co-operative mechanism whereby we *ask* a thread

to stop. The implementation code for a thread must be able to respond to such a request and perform the appropriate clean-up.

Unfortunately, in traditional (non-OSGi) Java development, these issues are rarely considered. Often threads are simply allowed to run to completion, or in some cases a thread is created during start-up of the application and assumed to run for as long as the Java process itself is running. For example a web server application would create a thread for accepting socket connections, and that thread would run until the server is shutdown. In these scenarios, there is no need to handle termination requests: the thread will be automatically stopped when Java stops.

In OSGi we don't have this luxury. Our bundle can be shutdown by a call to the `stop` method of its activator, and when that happens we *must* cleanup all artefacts that have been created on behalf of our bundle, including threads, sockets and so on. We cannot rely on our threads being forcibly ended by the termination of the JVM.

So how do we ask a thread to stop? Sadly Java doesn't even have a single consistent way to do this. There are a number of techniques, but we must choose the correct one based on what the thread is doing when we wish it to stop. Therefore in most cases we must tailor the termination code to the thread implementation code.

The code we saw in Chapter 2 (Listing 2.7) used the *interruption* mechanism, which is a simple boolean status that can be set on a thread by calling the `interrupt` method. It is possible for a thread to explicitly check its interruption status by calling `Thread.interrupted`, and also certain library methods (like `Thread.sleep` and `Object.wait`) are aware of their calling thread's interruption status and will exit immediately with an `InterruptedException` if the thread is interrupted while they are running.

Unfortunately, not all blocking library methods respond to interruption. For example most blocking I/O methods in `java.io` package simply ignore the interruption status; these methods continue to block, sometimes indefinitely, even after the thread has been interrupted. To wake them up we need some knowledge of what is blocking and why.

Listing 6.20 shows a server thread that accepts client connections. When a client connects, the server immediately sends the message "Hello, World!" to the client and then closes the connection to that client.

The activator here looks identical to `HeartbeatActivator` from Chapter 2: to stop the thread, we call its `interrupt` method. However in this case we have overridden the normal `interrupt` method with our own implementation that closes the server socket in addition to setting the interruption status. If the thread is currently blocking in the `ServerSocket.accept` method, closing the socket will cause it to exit immediately with an `IOException`. This pattern

Listing 6.20 Server Activator

```
1 package org.osgi.book.concurrency;
2
3 import java.io.IOException;
4 import java.io.PrintStream;
5 import java.net.ServerSocket;
6 import java.net.Socket;
7
8 import org.osgi.framework.BundleActivator;
9 import org.osgi.framework.BundleContext;
10
11 public class ServerActivator implements BundleActivator {
12
13     private HelloServer serverThread = new HelloServer();
14
15     public void start(BundleContext context) throws Exception {
16         serverThread.start();
17     }
18
19     public void stop(BundleContext context) throws Exception {
20         serverThread.interrupt();
21     }
22 }
23
24 class HelloServer extends Thread {
25
26     private static final int PORT = 9876;
27
28     private volatile ServerSocket socket = null;
29
30     public void interrupt() {
31         super.interrupt();
32         try {
33             if (socket != null) {
34                 socket.close();
35             }
36         } catch (IOException e) {
37             // Ignore
38         }
39     }
40
41     public void run() {
42         try {
43             socket = new ServerSocket(PORT);
44
45             while (!Thread.currentThread().isInterrupted()) {
46                 System.out.println("Accepting connections...");
47                 Socket clientSock = socket.accept();
48                 System.out.println("Client connected.");
49                 PrintStream out = new PrintStream(clientSock
50                     .getOutputStream());
51                 out.println("Hello, World!");
52                 out.flush();
53                 out.close();
54             }
55         } catch (IOException e) {
56             System.out.println("Server thread terminated.");
57         }
58     }
59 }
60
61 }
```

works for most of the `java.io` library: to interrupt a blocked I/O method, we can close the underlying socket or stream that it is blocking on.

As a rule of thumb, when you see a blocking library method that does *not* declare `InterruptedException` in its `throws` clause, it generally means it does not respond to interruption, and you must find another way to force it awake.

Much more discussion of these issues can be read in [15].

6.8 Exercises

1. Write a `ServiceTracker` that tracks a `Mailbox` services and stores them in a `Map<String, List<Mailbox>>`. The map should be keyed on the location of the publishing bundle (available from `ServiceReference.getBundle().getLocation()`) and the values should be the list of all `Mailbox` services published by that bundle. Also provide a `getMailboxes(String)` method that returns a snapshot of the services for a given bundle location. Ensure that any `synchronized` blocks used are as short as possible.
2. See the service interface in Listing 6.21. Write an implementation of this service which creates a thread pool of the size specified in the `size` field and registers it as an `Executor` service. The first time this method is called it should simply create and register the thread pool. On subsequent calls it should create and register a new thread pool, and also unregister and shutdown the old thread pool. Assume that your service class takes a `BundleContext` in its constructor.

Listing 6.21 Exercise 2: Thread Pool Manager Interface

```
1 package org.osgi.book.concurrency;
2
3 public interface ThreadPoolManager {
4     void updateThreadPool(int size);
5 }
```

7 The Whiteboard Pattern and Event Admin

So far our Mailbox Reader application is extremely simple; we can only display the headers of messages that were available in a mailbox at the time that the table was displayed. To be even remotely useful, the reader would also need to notify us when new messages arrive... otherwise we would have to restart the application to check for new messages!

The mechanism for new message notification variable across different types of mailboxes. With some communication protocols — such as POP3 email or RSS over HTTP — we must poll the server occasionally to ask if there are any new messages. Under other protocols such as SMS and IMAP, the server is able to “push” messages to our client as soon as they are available. Naturally the details of the notification mechanism need to be hidden inside the mailbox service implementation rather than exposed to clients of the service, such as our Reader GUI.

7.1 The Classic Observer Pattern

Most Java programmers will recognise the above as a problem that can be solved with the Observer pattern, also sometimes called the Listener pattern. And they would be correct, however OSGi adds a twist to the standard Java approach.

The classic pattern as defined by the “Gang of Four” [19] involves two entities: an Observable, which is the source of events; and a Listener, which consumes those events. The listener must be registered with an observable in order to receive events from it. Then the observable would notify all registered listeners of each event as it occurred. This is shown in Figure 7.1.

To use this pattern in our Mailbox Reader application, we could define a listener interface named `MailboxListener`. We would also need to extend the `Mailbox` interface to make it act as an observable, by adding methods to register and unregister listeners. These two interfaces are shown together in Listing 7.1.




Figure 7.1: The Classic Observer Pattern

Listing 7.1 Mailbox Listener and Observable Mailbox Interfaces

```

1 package org.osgi.book.reader.api;

3 public interface MailboxListener {
4     void messagesArrived(Mailbox mailbox, long[] ids);
5 }

1 package org.osgi.book.reader.api;

3 public interface ObservableMailbox extends Mailbox {
4     void addMailboxListener(MailboxListener listener);
5     void removeMailboxListener(MailboxListener listener);
6 }
  
```

7.2 Problems with the Observer Pattern

The observer pattern is very familiar, especially in GUI programming, but it has a number of problems. Most of these problems are general but they are particularly bad when we translate the pattern into an OSGi context.

The first problem is keeping track of dynamic observable objects. In our example, the source of events is a service that — like all services — has a life-cycle and can come and go many times during the life of the application. But when a new mailbox service appears, no listeners will be registered with it. Therefore any listener needs to hold open a **ServiceTracker** so it can be notified when new mailbox services appear and register itself, and likewise unregister itself when the service is going away. We must also find a way to deal with “lost updates”, i.e. events that occur between the mailbox appearing and the listener(s) registering themselves with it. If there are many listeners then they will not all register at the same time, so some listeners may see events that others do not see.

A second problem is that, just as the mailboxes can come and go, so can the listeners. Each observable therefore needs to manage a changing set of listeners, and make sure it does not attempt to deliver events to listeners that have disappeared. Since the registering and unregistering of listeners can occur in a different thread from the firing of an event, proper synchronization must

be used to avoid concurrency bugs.

There is a memory management problem here also: the observer pattern is one of the principal causes of memory leaks in Java applications. When a listener is registered with an observable, the internal implementation of the observable will typically add that listener to collection field. But now there is a strong reference to the listener object merely because it exists inside the collection, preventing it from being cleaned up by the garbage collector. Even if the listener is not useful any more in its original context, it will live on in memory until the observer dies.

Therefore it is very important to clean up listeners when they are no longer required. But this is very difficult to verify, and is often not done correctly. The problem is that it is very easy to detect a problem with the set-up phase — if done incorrectly, the application simply does not work — but problems with cleaning up do not directly break the application, instead causing “out of memory” conditions much later on.

7.3 Fixing the Observer Pattern

None of the problems above are insurmountable, of course. By writing both our listeners and observers carefully and correctly we can avoid all concurrency bugs and memory leaks. The remaining problem, however, is that there is quite a lot of complex code to write each time we want to use the observer pattern, and it is not good practice to repeat such code many times. Therefore we need to look for a way to centralise that code in a single reusable component or library, so that we can go back to writing “dumb” observers and listeners.

One possible solution is to create some kind of event broker that sits between the event sources and the event listeners. The listeners could register themselves against the broker rather than directly with each observable, and the event sources would simply publish events to the broker and allow it to deliver those events to all registered listeners. Figure 7.2 shows the general idea.

How does this simplify matters for the listeners? Instead of having to track each observer and register and unregister themselves each time an observer appears and disappears, the listeners would simply need to register themselves *once* with the broker and unregister once when no longer needed. We can assume the broker will have a much longer life than the individual observers or listeners, so we will not have to worry about the broker itself going away.

And for the event sources? They now have a much easier job. Rather than managing a variable size collection of listeners, they can simply publish their events to a single place. In fact we no longer need special observable interfaces such as `ObservableMailbox` from Listing 7.1. Now, *any* object can publish

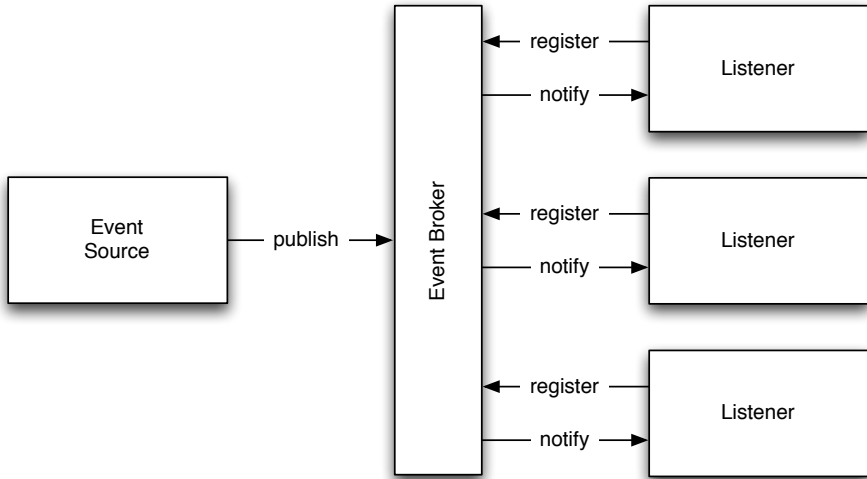


Figure 7.2: An Event Broker

events to all of the current listeners for that event type simply by calling the event broker. So, both event sources and listeners can now be written with dumb code.

An alternative to the event broker is a “listener directory”, as in Figure 7.3. This would take responsibility for managing the listeners for a particular kind of event, but it would not actually deliver the events itself. Instead it would provide a snapshot of the current set of listeners, enabling the event source to deliver the events. This is slightly more flexible, since the event source could choose to follow a custom delivery strategy. For example it may wish to deliver events to multiple listeners in parallel using a thread pool.

Of course, either the event broker or the listener directory implementations will be quite complex and must be written with great care, although fortunately they need only be written once. However, in OSGi we are lucky because an implementation of both these patterns already exists! In fact the listener directory is simply the Service Registry that we are already familiar with.

7.4 Using the Whiteboard Pattern

The listener directory pattern we have just described is more poetically known as the Whiteboard Pattern. We will see why this name fits in a moment, but let’s look at how we can use the Service Registry as a directory of listeners.

To register a listener, we simply take that listener object and publish it to the

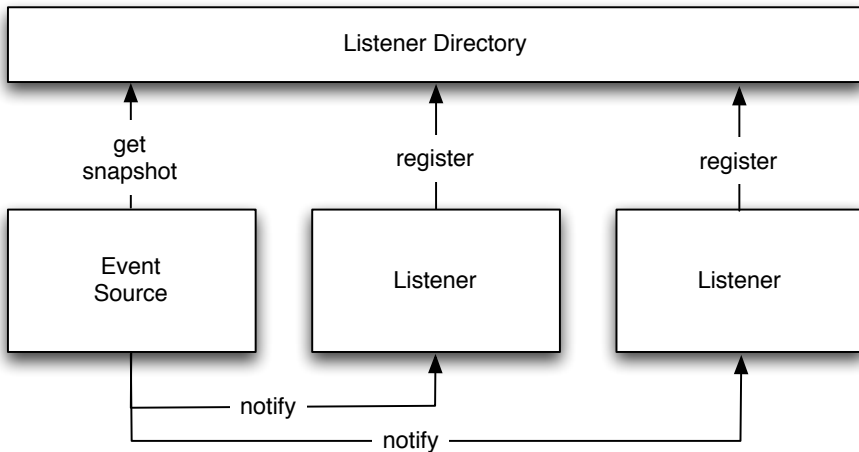


Figure 7.3: A Listener Directory

service registry under its specific `Listener` interface. For example, Listing 7.2 shows a bundle activator that registers an implementation of the `MailboxListener` interface. This is literally *all we have to do* on the listener side.

Listing 7.2 Registering a Mailbox Listener

```

1 package org.osgi.book.reader.gui;
2
3 import org.osgi.book.reader.api.MailboxListener;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6
7 public class MailboxListenerActivator implements BundleActivator {
8
9     public void start(BundleContext context) throws Exception {
10         MailboxListener listener = new MyMailboxListener();
11         context.registerService(MailboxListener.class.getName(),
12                                listener, null);
13     }
14
15     public void stop(BundleContext context) throws Exception {
16         // No need to explicitly unregister the service
17     }
18 }
  
```

This simplicity explains the name “Whiteboard Pattern”. We do not need to actively seek out event sources and explicitly register against them, instead we merely declare the existence of the listener and wait against event sources to find it. This is similar to writing your name on a whiteboard in order to declare your interest in joining an activity, rather than finding the organiser of that activity and registering your interest directly.

Then how does an event source interact with the whiteboard pattern? Well, it must do more than just registering a service, but not a lot more. Typically it would use a `ServiceTracker` and call the `getServices` method to get a snapshot of all the currently registered listeners, which it can then iterate over. An example based on the `MailboxListener` interface is shown in Listing 7.3. Here we subclass `ServiceTracker` and offer a `fireMessagesArrived` method that performs the work.

Listing 7.3 Mailbox Listener Tracker

```
1 package org.osgi.book.reader.asyncmailbox;
2
3 import org.osgi.book.reader.api.Mailbox;
4 import org.osgi.book.reader.api.MailboxListener;
5 import org.osgi.framework.BundleContext;
6 import org.osgi.util.tracker.ServiceTracker;
7
8 public class MailboxListenerTracker extends ServiceTracker {
9
10     public MailboxListenerTracker(BundleContext context) {
11         super(context, MailboxListener.class.getName(), null);
12     }
13
14     public void fireMessagesArrived(Mailbox mailbox, long[] ids) {
15         Object[] services = getServices();
16         if(services != null) {
17             for (Object service : services) {
18                 ((MailboxListener) service).messagesArrived(mailbox, ids);
19             }
20         }
21     }
22 }
```

This `MailboxListenerTracker` class is an obvious candidate for reuse: all mailbox implementations that have the capability to receive messages asynchronously will probably want to use this class or something very similar. So, we could pull it out into a shared library or even supply it as a helper class in the base Mailbox API.

However, we can do better than this. There is an even more general pattern at work in this class, which readers of “Gang of Four” will recognise as the Visitor pattern. Given this observation we can write a general purpose helper class that can be reused every time we implement the whiteboard pattern. This is shown in Listing 7.4, in which we define a `WhiteboardHelper` class that takes instances of a `Visitor` interface. To use the helper in our Mailbox example we can simply pass in an instance of `Visitor<MailboxListener>` that makes a call to the `messagesArrived` method. Java 5 generics are used to make the pattern type-safe but it can also work under Java 1.4 using casts.

Listing 7.4 Visitor Interface and Whiteboard Helper Class

```

1 package org.osgi.book.utils;

3 public interface Visitor<T> {
4     void visit(T object);
5 }

1 package org.osgi.book.utils;

3 import org.osgi.framework.BundleContext;
4 import org.osgi.util.tracker.ServiceTracker;

6 public class WhiteboardHelper<T> extends ServiceTracker {

8     public WhiteboardHelper(BundleContext context, Class<T> svcClass) {
9         super(context, svcClass.getName(), null);
10    }

12    public void accept(Visitor<? super T> visitor) {
13        Object[] services = getServices();
14        if (services != null) {
15            for (Object serviceObj : services) {
16                @SuppressWarnings("unchecked")
17                T service = (T) serviceObj;

19                visitor.visit(service);
20            }
21        }
22    }
23 }

```

7.4.1 Registering the Listener

Now let's have a look at how to introduce the whiteboard pattern into our example Mailbox Reader application, starting with the listener side. In the Mailbox Reader application, each table displaying the contents of a mailbox is likely to be interested in receiving notifications of new messages. In a real-world application, other parts of the GUI are also likely to be interested in these notifications — for example we might want to show an indicator icon in the “System Tray” that pops up messages when a new message arrives. The advantage of the whiteboard approach is that components of the GUI can independently register `MailboxListener` services and receive notifications while remaining oblivious to the existence of other listeners. This enhances the modularity of the application.

In this example we will implement the functionality to update the mailbox table display when new messages arrive. We first need to decide which object will implement the `MailboxListener` interface: it should be one which is close to the model and is able to respond to updates to that model and notify the Swing GUI framework. The obvious choice here is the `MailboxTableModel` class, which inherits several such notification methods from its superclass `AbstractTableModel`. The one we will want to call when new messages arrive

is `fireTableRowsInserted`. Listing 7.5 shows only the new code we need to add to the class.

Listing 7.5 Adding the MailboxListener Interface to MailboxTableModel

```

1 package org.osgi.book.reader.gui;
2
3 // Imports omitted ...
4
5 public class MailboxTableModel extends AbstractTableModel implements
6     MailboxListener {
7
8     // Previously defined methods omitted ...
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64     public void messagesArrived(Mailbox mailbox, long[] ids) {
65         if (mailbox != this.mailbox) {
66             // Ignore events for other mailboxes
67             return;
68         }
69
70         List<Message> newMessages;
71         try {
72             newMessages = Arrays.asList(mailbox.getMessages(ids));
73         } catch (MailboxException e) {
74             newMessages = Collections.emptyList();
75         }
76
77         final int firstNew, lastNew;
78         synchronized (this) {
79             firstNew = messages.size(); // Index of the first new row
80             messages.addAll(newMessages);
81             lastNew = messages.size() - 1; // Index of the last new row
82         }
83
84         SwingUtilities.invokeLater(new Runnable() {
85             public void run() {
86                 fireTableRowsInserted(firstNew, lastNew);
87             }
88         });
89     }
90 }

```

The next problem is how we register the table model as a service under the `MailboxListener` interface. It makes sense to do this inside `MailboxPanel` because that is the class in which the table model is created and held. We simply add `registerListener` and `unregisterListener` methods to the class. The former needs to take the `BundleContext` as a parameter, but the latter takes no parameter; it simply calls `unregister` on the `ServiceRegistration` object that was created in the `registerListener` method.

The final modification we must make is to include calls to `registerListener` and `unregisterListener` from the `addingService` and `removedService` methods of `ScannerMailboxTracker`. We will not repeat the entire code of the tracker here, since we must only insert two lines:

- In `addingService`, after construction of the new `MailboxPanel` object

Listing 7.6 Mailbox Panel, with MailboxListener Registration

```
1 package org.osgi.book.reader.gui;
3 // Imports omitted...
5 public class MailboxPanel extends JPanel {
7     private final MailboxTableModel tableModel;
8     private volatile ServiceRegistration svcReg;
10    // Existing constructor omitted...
26    public void registerListener(BundleContext context) {
27        svcReg = context.registerService(
28            MailboxListener.class.getName(), tableModel, null);
29    }
31    public void unregisterListener() {
32        if(svcReg == null) throw new IllegalStateException();
33        svcReg.unregister();
34    }
35 }
```

during the `Callable.call` method, we insert:

```
panel.registerListener(context);
```

- In `removedService`, before removing the `MailboxPanel` from the tabbed pane during the `Runnable.run` method, we insert:

```
panel.unregisterListener();
```

7.4.2 Sending Events

Now we look at implementing the event source side of the whiteboard pattern: a mailbox that sends events when new messages arrive.

However, first we need to implement a mailbox which actually does receive new messages! Our previous sample mailbox was fixed in size, with a hard-coded list of initial messages, so now we need a mailbox that can grow. A realistic implementation would involve network sockets and polling and other complexity, so again we will keep things simple and implement a mailbox that grows by one message every five seconds. To achieve this we drive a timer thread from our bundle activator, in a very similar way to the `HeartbeatActivator` example from Chapter 2.

Listing 7.7 shows the activator with the timer thread. We assume that we have a mailbox implementation class `GrowableMailbox` that has an `addMessage` method, allowing messages to be added programmatically. `GrowableMailbox` also takes a `WhiteboardHelper` in its constructor, so we create this in the activator's `start` method before constructing the mailbox.

Listing 7.7 Growable Mailbox Activator and Timer Thread

```
1 package org.osgi.book.reader.asyncmailbox;
2
3 import java.util.Date;
4 import java.util.Properties;
5
6 import org.osgi.book.reader.api.Mailbox;
7 import org.osgi.book.reader.api.MailboxListener;
8 import org.osgi.book.utils.WhiteboardHelper;
9 import org.osgi.framework.BundleActivator;
10 import org.osgi.framework.BundleContext;
11 import org.osgi.framework.ServiceRegistration;
12
13 public class GrowableMailboxActivator implements BundleActivator {
14
15     private WhiteboardHelper<MailboxListener> whiteboard;
16     private GrowableMailbox mailbox;
17     private Thread messageAdderThread;
18     private ServiceRegistration svcReg;
19
20     public void start(BundleContext context) throws Exception {
21         whiteboard = new WhiteboardHelper<MailboxListener>(context,
22                                                         MailboxListener.class);
23         whiteboard.open();
24
25         mailbox = new GrowableMailbox(whiteboard);
26         messageAdderThread = new Thread(new MessageAdder());
27         messageAdderThread.start();
28
29         Properties props = new Properties();
30         props.put(Mailbox.NAME_PROPERTY, "growing");
31         svcReg = context.registerService(Mailbox.class.getName(),
32                                         mailbox, props);
33     }
34
35     public void stop(BundleContext context) throws Exception {
36         svcReg.unregister();
37         messageAdderThread.interrupt();
38         whiteboard.close();
39     }
40
41     private class MessageAdder implements Runnable {
42         public void run() {
43             try {
44                 while (!Thread.currentThread().isInterrupted()) {
45                     Thread.sleep(5000);
46                     mailbox.addMessage("Message added at " + new Date(),
47                                       "Hello again");
48                 }
49             } catch (InterruptedException e) {
50                 // Exit quietly
51             }
52         }
53     }
54 }
```

To implement the “growable” mailbox itself, we subclass it from `FixedMailbox` to inherit the internal storage of messages. We just need to add a protected `addMessage` method to be called from the timer thread. This is shown in Listing 7.8. Note the `synchronized` block surrounding access to the internal `messages` field. The superclass must also synchronise accesses to that field, which is the reason why all the methods of `FixedMailbox` were declared `synchronized` (this was left unexplained when that class was first introduced in Chapter 3).

Here is where we see the power of the whiteboard pattern at work. Having added a new message to its internal data structure, `GrowableMailbox` creates a visitor object that calls `messagesArrived` against each listener that it visits.

Listing 7.8 Growable Mailbox

```
1 package org.osgi.book.reader.asyncmailbox;
2
3 import org.osgi.book.reader.api.Mailbox;
4 import org.osgi.book.reader.api.MailboxListener;
5 import org.osgi.book.reader.api.Message;
6 import org.osgi.book.reader.fixedmailbox.FixedMailbox;
7 import org.osgi.book.reader.fixedmailbox.StringMessage;
8 import org.osgi.book.utils.Visitor;
9 import org.osgi.book.utils.WhiteboardHelper;
10
11 public class GrowableMailbox extends FixedMailbox {
12
13     private final WhiteboardHelper<MailboxListener> whiteboard;
14
15     public GrowableMailbox(WhiteboardHelper<MailboxListener> wb) {
16         this.whiteboard = wb;
17     }
18
19     protected void addMessage(String subject, String text) {
20         final int newMessageId;
21
22         synchronized (this) {
23             newMessageId = messages.size();
24             Message newMessage = new StringMessage(newMessageId,
25                                                     subject, text);
26             messages.add(newMessage);
27         }
28
29         final long[] newMessageIds = new long[] { newMessageId };
30         final Mailbox source = this;
31         Visitor<MailboxListener> v = new Visitor<MailboxListener>() {
32             public void visit(MailboxListener l) {
33                 l.messagesArrived(source, newMessageIds);
34             }
35         };
36         whiteboard.accept(v);
37     }
38 }
```

7.5 Event Admin

The whiteboard pattern works well in many situations, such as the one described above, however in some cases it can add overhead that is undesirable.

For example, the event sources must implement all of the code to iterate over the listener services. The `WhiteboardHelper` is rather simplistic: on the first error it stops delivering messages to other listeners that might be registered, and if any listener takes a long time to process an event it will hold up the execution of the event source (i.e., the `addMessage` method of `GrowableMailbox` will not return until all listeners have processed the newly added message). Also, the event sources are still somewhat coupled to the consumers, since they can only deliver events to implementers of a specific interface. Therefore both the event source and the consumer must depend on the same version of that interface.

Sometimes we would like to create an event source that sends many fine-grained events to many different listeners. For example, a financial trading application may need to deliver quotes (e.g., stock quotes, foreign exchange rates, etc) to many consumers of that data. The whiteboard pattern does not scale well in that scenario. Instead we should look at the other pattern described in Section 7.3: the *Event Broker*. OSGi provides us with an implementation of this pattern called the Event Admin Service, which is a service specified in the OSGi Service Compendium. Event Admin implements a form of *publish/subscribe*, which is popular in many message-based systems.

The operation of Event Admin is slightly different from the diagram in Figure 7.2. In that diagram, the listeners registered themselves directly with the event broker, but Event Admin actually uses the whiteboard pattern: listeners register as services under the `EventHandler` interface, and the “broker” tracks them. The diagram for Event Admin is therefore a combination of both Figures 7.2 and 7.3, and is shown in Figure 7.4.

7.5.1 Sending Events

Because it is based on the whiteboard pattern, using Event Admin from the consumer end looks very similar to what we have already seen. So we will first look at the event source end, which does look substantially different. Rather than tracking each listener service and iterating over them each time an event must be sent, when using Event Admin our sources simply need to make a single method call to the Event Admin service.

For our example we will model a financial trading application which continually receives “market data”, i.e., prices of various tradable assets. Again, to avoid the incidental complexity of doing this for real we will simulate it by generating

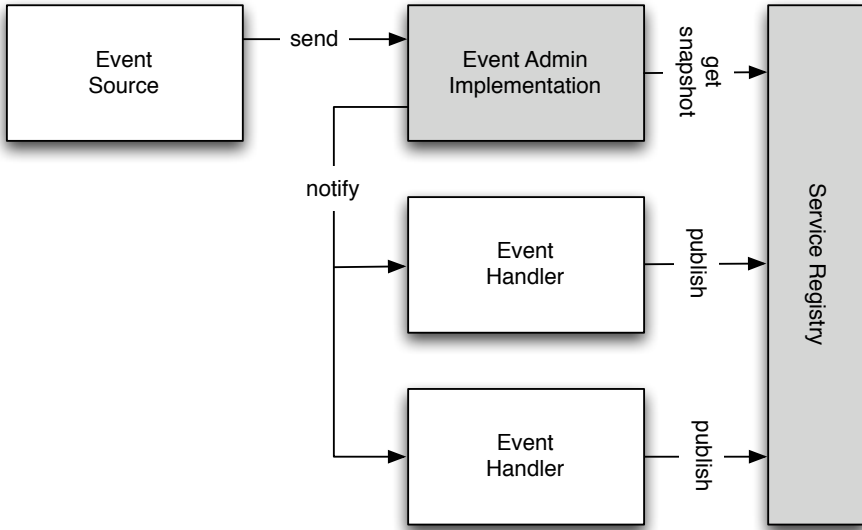


Figure 7.4: The Event Admin Service

random data on a timer thread. Listing 7.9 shows the implementation of a `Runnable` that does just this: it sends stock quotes for the symbol `MSFT` (Microsoft Corporation) with a starting price of 25 and randomly adjusting upwards or downwards every 100 milliseconds.

This class requires an instance of `EventAdmin` to be passed in its constructor. A good way to provide this is by following the pattern used in the `LogTracker` class from Chapter 4. That is, we subclass `ServiceTracker` and implement the `EventAdmin` interface with delegating calls to the result of calling `getService`. The result is `EventAdminTracker` shown in Listing 7.10

Listing 7.11 shows the activator which manages the `EventAdminTracker` and the timer thread.

7.5.2 The Event Object

The `Event` object what we pass to the Event Admin service is an immutable object that consists of a topic and an arbitrary set of properties.

The topic defines the logical type of the event, and its primary purpose is to act as a filter to determine which handlers should receive which events. It consists of a sequence of tokens separated by slashes, which serve to form a hierarchy. This allows consumers to filter out at any level of the hierarchy, so in the preceding example a consumer could choose to receive: all prices;

Listing 7.9 Random Price Generator

```
1 package org.osgi.book.trading.feeds;
2
3 import java.util.Properties;
4 import java.util.Random;
5
6 import org.osgi.service.event.Event;
7 import org.osgi.service.event.EventAdmin;
8
9 public class RandomPriceFeed implements Runnable {
10
11     private static final String TOPIC = "PRICES/STOCKS/NASDAQ/MSFT";
12
13     private final EventAdmin eventAdmin;
14
15     public RandomPriceFeed(EventAdmin eventAdmin) {
16         this.eventAdmin = eventAdmin;
17     }
18
19     public void run() {
20         double price = 25;
21         Random random = new Random();
22
23         try {
24             while (!Thread.currentThread().isInterrupted()) {
25                 // Create and send the event
26                 Properties props = new Properties();
27                 props.put("symbol", "MSFT");
28                 props.put("time", System.currentTimeMillis());
29                 props.put("price", price);
30                 eventAdmin.sendEvent(new Event(TOPIC, props));
31
32                 // Sleep 100ms
33                 Thread.sleep(100);
34
35                 // Randomly adjust price by upto 1.0, + or -
36                 double nextPrice = random.nextBoolean()
37                     ? price + random.nextDouble()
38                     : price - random.nextDouble();
39                 price = Math.max(0, nextPrice);
40             }
41         } catch (InterruptedException e) {
42             // Exit quietly
43         }
44     }
45 }
```

Listing 7.10 Event Admin Tracker

```
1 package org.osgi.book.utils;
2
3 import org.osgi.framework.BundleContext;
4 import org.osgi.service.event.Event;
5 import org.osgi.service.event.EventAdmin;
6 import org.osgi.util.tracker.ServiceTracker;
7
8 public class EventAdminTracker extends ServiceTracker
9         implements EventAdmin {
10
11     public EventAdminTracker(BundleContext context) {
12         super(context, EventAdmin.class.getName(), null);
13     }
14
15     public void postEvent(Event event) {
16         EventAdmin evtAdmin = (EventAdmin) getService();
17         if(evtAdmin != null) evtAdmin.postEvent(event);
18     }
19
20     public void sendEvent(Event event) {
21         EventAdmin evtAdmin = (EventAdmin) getService();
22         if(evtAdmin != null) evtAdmin.sendEvent(event);
23     }
24 }
```

Listing 7.11 Bundle Activator for the Random Price Generator

```
1 package org.osgi.book.trading.feeds;
2
3 import org.osgi.book.utils.EventAdminTracker;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6
7 public class RandomPriceFeedActivator implements BundleActivator {
8
9     private EventAdminTracker evtAdmTracker;
10    private Thread thread;
11
12    public void start(BundleContext context) throws Exception {
13        evtAdmTracker = new EventAdminTracker(context);
14        evtAdmTracker.open();
15
16        thread = new Thread(new RandomPriceFeed(evtAdmTracker));
17        thread.start();
18    }
19
20    public void stop(BundleContext context) throws Exception {
21        thread.interrupt();
22        evtAdmTracker.close();
23    }
24 }
```

all *stock* prices; just prices for stocks traded on NASDAQ; or just prices for Microsoft Corporation.

Data describing the actual event should appear in the properties, which is actually an instance of `Dictionary`. So in this example the time and price of the quote are included, as well as the stock symbol. Note that in general, data describing the event should *not* be added to the topic. However a field that is likely to be used as a primary filter (such as the stock symbol in this example) could appear in both the topic and the properties. Only one such field should be used this way, though.

The immutability of the `Event` class is a very important feature. If it were mutable, then any handler could change the content of an event, and therefore other handlers receiving it subsequently would see the altered content instead of the original event created by the sender. To support the immutability of `Event`, the specification for Event Admin states that only instances of `String` and the eight primitive types (i.e., `int`, `float`, `double` etc). or their Object wrappers (`Integer`, `Float`, `Double`) may be added to the set of properties for an event.

This may seem rather restrictive, and it is tempting to ignore the rule and add more complex objects directly to the event's `Dictionary`. For example we may have a `Quote` class implemented as a standard `JavaBean` with getters and setters. There is nothing to stop us adding such an object¹ but it is not a good idea to do so. If `Quote` is mutable then the event itself will be effectively mutable. More seriously, we would introduce some coupling between the event sources and consumers: only consumers that import the `Quote` class (and furthermore, the same *version* of it) will be able to conveniently access the contents of the event. Other consumers could use Java reflection to look at the quote object but the code to do this is far more cumbersome and less efficient than simply accessing a `Dictionary`.

Another reason for using only simple values in the event properties is because we can then apply a secondary filter to events based on those values. We will see an example in the next section.

7.5.3 Receiving Events

Since we have already implemented the whiteboard pattern, we already know how to implement an event handler for use with Event Admin: simply register services under the `EventHandler` interface.

However, we must supply at least a topic declaration as a service property along with the registration. This tells the Event Admin broker which topics we are interested in, and we can either specify the topic in full or use a wildcard in

¹At compile time; some Event Admin implementations *may* enforce the rule at runtime.

order to receive messages on all topics beneath a certain level in the hierarchy. For example specifying `PRICES/STOCKS/*` will give us all stock prices and `PRICES/STOCKS/NYSE/*` will give us just prices published on the New York Stock Exchange. We can also specify `*` alone to receive *all* events on all topics. Note that we cannot place a `*` in the middle of the topic string, i.e., `PRICES/*/NYSE...` is not allowed.

The name for the topic property is “`event.topics`”, but from Java code we tend to use a static constant, `EventConstants.EVENT_TOPIC`. This property is mandatory: `EventHandler` services that do not specify a topic will not receive any events.

Another property we can set is “`event.filter`” or `EventConstants.EVENT_FILTER`. This property is optional but it allows us to apply an additional filter on the contents of the event properties, using an LDAP-style query.

Listing 7.12 shows an example of registering an `EventHandler`, which prints stock quotes to the console. It uses a filter to print only prices greater than or equal to 20.

7.5.4 Running the Example

As Event Admin is a compendium service, we need to compile against the compendium API JAR, `osgi.cmpn.jar`. That JAR also needs to be installed in Felix as a bundle to provide the API to our bundles at runtime.

To get the example working we will also need to install an implementation of the Event Admin service. Felix supplies one (as do all the other OSGi frameworks) but it is not included in the default download. However we can install and start it directly from the Felix console as follows (NB: do not include the line break in the middle of the URL):

```
-> install http://www.apache.org/dist/felix/org.apache.felix
    .eventadmin-1.0.0.jar
Bundle ID: 26
-> start 26
->
```

Now we can verify that the Event Admin service is running by typing the `services` command. We should see the following somewhere in the output:

```
Apache Felix EventAdmin (26) provides:
-----
org.osgi.service.event.EventAdmin
```

To build the sender and receiver bundles, use the two `bnd` descriptors shown in Listing 7.13.

After building we can install and start the bundles:

Listing 7.12 Stock Ticker Activator

```

1 package org.osgi.tutorial;
2
3 import java.util.Properties;
4
5 import org.osgi.framework.BundleActivator;
6 import org.osgi.framework.BundleContext;
7 import org.osgi.service.event.Event;
8 import org.osgi.service.event.EventConstants;
9 import org.osgi.service.event.EventHandler;
10
11 public class StockTickerActivator implements BundleActivator {
12
13     private static final String STOCKS_TOPIC = "PRICES/STOCKS/*";
14
15     public void start(BundleContext context) throws Exception {
16         EventHandler handler = new EventHandler() {
17             public void handleEvent(Event event) {
18                 String symbol = (String) event.getProperty("symbol");
19                 Double price = (Double) event.getProperty("price");
20
21                 System.out.println("The price of " + symbol
22                                     + " is now " + price);
23             }
24         };
25
26         Properties props = new Properties();
27         props.put(EventConstants.EVENT_TOPIC, STOCKS_TOPIC);
28         props.put(EventConstants.EVENT_FILTER, "(price>=20)");
29         context.registerService(EventHandler.class.getName(),
30                                 handler, props);
31     }
32
33     public void stop(BundleContext context) throws Exception {
34     }
35
36 }

```

Listing 7.13 Bnd Descriptors for the Random Price Feed and Stock Ticker

```

1 # random_feed.bnd
2 Private-Package: org.osgi.book.trading.feeds,org.osgi.book.utils
3 Bundle-Activator: org.osgi.book.trading.feeds.RandomPriceFeedActivator
4
5 # ticker.bnd
6 Private-Package: org.osgi.tutorial
7 Bundle-Activator: org.osgi.tutorial.StockTickerActivator

```

```
-> install file:random_feed.jar
Bundle ID:27
-> install file:ticker.jar
Bundle ID: 28
-> start 27
-> start 28
The price of MSFT is now 25.1945254744465635
The price of MSFT is now 24.547478302312108
The price of MSFT is now 25.173540992244572
The price of MSFT is now 25.906669217574922
The price of MSFT is now 25.41915022996729
The price of MSFT is now 26.04457130652444
The price of MSFT is now 26.29259735036186
The price of MSFT is now 25.64594680159028
The price of MSFT is now 26.279434082391102
The price of MSFT is now 27.012694572863026
...
```

7.5.5 Synchronous versus Asynchronous Delivery

Event Admin includes a much more sophisticated event delivery mechanism than the simplistic `WhiteboardHelper` class. One of the most important differences is that Event Admin can optionally deliver events to handlers *asynchronously*.

Most examples of the whiteboard pattern — including our `WhiteboardHelper` class — deliver events synchronously, meaning that each listener is called in turn from the event source thread, and therefore the event source cannot continue any other processing until all of the listeners have finished processing the event, one by one.

Event Admin does support synchronous processing, and that is what we used in the previous example, however it also supports asynchronous processing, which means that the call to the broker will return immediately, allowing the event source to continue with other processing. The events will be delivered to the listeners in one or more threads created by Event Admin for this purpose. To use asynchronous processing we simply call the `postEvent` method of `EventAdmin` rather than `sendEvent`. Nothing else (in our code) needs to change.

Using `postEvent` to request asynchronous delivery makes Event Admin behave much more like a true event broker, albeit an in-JVM one. For the event source it is “fire and forget”, allowing us to quickly post an event or message and let Event Admin worry about delivery to the end consumers. Therefore you should in general use `postEvent` as a preference, unless you have a particular need to wait until an event has been delivered to all consumers, in which case you should use `sendEvent`.

Note that using `sendEvent` does *not* guarantee that the `handleEvent` method of all handlers will actually be called in the same thread that the event source used to call `sendEvent`. Event Admin may choose to use several threads to

deliver the event to several handlers concurrently, even when we are using synchronous delivery — the guarantee given by `sendEvent` is merely that it will return to the caller only after the event has been fully delivered to all handlers. As a consequence, we cannot make any assumption about the thread on which handlers will receive events, even if we believe we know which thread is sending them.

Another reason to favour `postEvent` over `sendEvent` is we can induce deadlocks if we hold a lock when calling `sendEvent`, as described at length in Chapter 6. However it is practically impossible to cause deadlock by holding a lock when calling `postEvent`.

7.5.6 Ordered Delivery

Whether we are using synchronous or asynchronous delivery, Event Admin promises to deliver events to each handler in the same order in which they arrived at the broker. That is, if a single thread sends or posts events *A*, *B* and *C* in that order to Event Admin, then each handler will see events *A*, *B* and *C* arriving in that order. However if multiple threads are sending or posting events to Event Admin, the order in which those events arrive at the broker depends on low level timing factors and so the delivery order to handlers is not guaranteed. That is, if a second thread sends events *D*, *E* and *F* in that order, then each handler may see the order *ABCDEF* or perhaps *DEFABC*, *ADBECE*, *DAEBFC* or some other interleaving. But they will not see *CABDEF*, *ADCFBE* etc., since these examples violate the internal ordering of messages within each thread.

7.5.7 Reliable Delivery

Event Admin attempts to make the delivery of events reliable when faced with misbehaving handlers. If a handler throws an exception then Event Admin will catch it and log it the OSGi Log Service, if it is available. Then it will continue delivering the event to other handlers. Note it does not attempt to catch Errors such as `OutOfMemoryError`, `LinkageError` etc. Also because the `EventHandler.handleEvent` method does not declare any checked exceptions in a `throws` clause, we can only throw subclasses of `RuntimeException` from our handlers.

Some implementations of Event Admin may also attempt to detect handlers that have stalled, for example in an infinite loop or deadlock. They may choose “blacklist” particular misbehaving handlers so they no longer receive any events. However, this feature is optional and not all Event Admin implementations support it.

7.6 Exercises

1. The notification mechanism shown in Section 7.4 is somewhat inefficient because all mailbox events are sent to all registered mailbox listeners, but many of the listeners (e.g., the message tables) are only interested in the events for a *single* mailbox. Extend this code to implement filtering based on mailbox name. Each listener should be able to listen to just one mailbox or all mailboxes. Hint: `WhiteboardHelper` will need to be extended to accept a filter expression.

8 The Extender Model

We saw in Chapter 4 how services can be used to extend the functionality of an application at runtime with new Java classes. Sometimes though we want extensibility of another sort: the ability to add artifacts other than executable code.

A good example is the help system of our Mailbox Reader sample application. We are able to extend the functionality of the application by plugging in new bundles, such as bundles providing new mailbox types, which are naturally implemented with services. But we also need to provide documentation for these new features. Therefore our help system needs to be extensible, too.

Let's assume that help documents will be in HTML format. Documentation is usually static, so a plain HTML file will suffice, i.e. the content of the HTML document does not need to be generated on-the-fly. Now, we *could* use a service to register the existence of a file, by declaring an interface as in Listing 8.1.

Listing 8.1 Help Provider Service Interface

```
1 public interface HelpDocumentProvider {
2     /**
3      * Return the URL of an HTML file.
4      */
5     URL getHelpDocumentURL ();
6 }
```

But this would be a very cumbersome way to accomplish such a simple task. Any bundle wishing to provide help would have to implement the interface, and also implement `BundleActivator` in order to instantiate the class and register it as a service... all this just to provide a simple URL!

It would be a lot easier if we didn't need to "register" the document at all, but instead the help system simply found it inside our bundle and processed it automatically. Perhaps we could give a hint to the help system about which HTML files are intended for its use, e.g. by placing those files in a `help` subdirectory of the bundle.

It turns out this is quite easy to achieve, because bundles can see the contents of other bundles. So we can write a bundle that scans other bundles, looking

for help documentation and adding it to the central index. We call this kind of bundle an *extender*.

An extender bundle is one which scans the contents and/or headers of other bundles and performs some action on their behalf. This is a very useful and common pattern in OSGi.

8.1 Looking for Bundle Entries

Let's take a look at the code required to look at the contents of another bundle. We have two tools for doing this, both of them methods on the `Bundle` interface. The first is `getEntry`, which takes a path string such as `help/index.html` and returns a URL reference to that entry in the bundle if it exists; the contents of the entry can be read from the URL by calling `openStream` on it. The second is `getEntryPaths` which takes a prefix string and returns an enumeration of all the bundle entry paths starting with that prefix.

The method in Listing 8.2 will scan a bundle and return a list of URLs pointing to the HTML documents we are interested in, i.e. those that have been placed under the `help` subdirectory. Notice, again, the necessity of doing a null-check on the result of the query method: OSGi never returns an empty array or collection type. However our method *does* return an empty List of URLs when no matches are found.

Listing 8.2 Scanning a Bundle for Help Documents

```
17 private List<URL> scanForHelpDocs(Bundle bundle) {
18     List<URL> result;
19     Enumeration<?> entries = bundle.getEntryPaths("help");
20     if (entries != null) {
21         result = new ArrayList<URL>();
22         while (entries.hasMoreElements()) {
23             String entry = (String) entries.nextElement();
24             if (entry.endsWith(".html")) {
25                 result.add(bundle.getEntry(entry));
26             }
27         }
28     } else {
29         result = Collections.emptyList();
30     }
31     return result;
32 }
```

This code works just fine but is a little limited. The problem is we need to take the information returned by `scanForHelpDocs` and use it to create an index page in our help system, but all we know about each document is its filename. Therefore all we could show in the help index would be a list of filenames, which is not likely to be very helpful.

We could remedy this by asking help providers to list all the HTML documents they provide explicitly and supply a title using a simple properties file as shown in Listing 8.3.

Listing 8.3 A Help Index File, `index.properties`

```
introduction=Introduction
first_steps=First Steps in OSGi
dependencies=Bundle Dependencies
intro_services=Introduction to Services
```

We can interpret this as follows: the file named `help/introduction.html` has the title “Introduction”; the file named `help/first_steps.html` has the title “First Steps in OSGi”; and so on.

Let’s assume that this properties file can be found at the location `help/index.properties`. The code in Listing 8.4 is an improved version of the scanning method, which now returns not just a list of document URLs but a list of URLs and titles:

Listing 8.4 Scanning a Bundle for Help Documents with Titles (1)

```
34 private List<Pair<URL, String>> scanForHelpDocsWithTitles(
35     Bundle bundle) throws IOException {
36     // Find the index file entry; exit if not found
37     URL indexEntry = bundle.getEntry("help/index.properties");
38     if (indexEntry == null) {
39         return Collections.emptyList();
40     }
41
42     // Load the index file as a Properties object
43     Properties indexProps = new Properties();
44     InputStream stream = null;
45     try {
46         stream = indexEntry.openStream();
47         indexProps.load(stream);
48     } finally {
49         if (stream != null)
50             stream.close();
51     }
52
53     // Iterate through the files
54     List<Pair<URL, String>> result =
55         new ArrayList<Pair<URL, String>>(indexProps.size());
56     Enumeration<?> names = indexProps.propertyNames();
57     while (names.hasMoreElements()) {
58         String name = (String) names.nextElement();
59         String title = indexProps.getProperty(name);
60
61         URL entry = bundle.getEntry("help/" + name + ".html");
62         if (entry != null) {
63             result.add(new Pair<URL, String>(entry, title));
64         }
65     }
66
67     return result;
68 }
```

Note that in order to represent a document URL together with a title, we define a simple `Pair` class, the definition of which is shown in Listing 8.5. This uses Java 5 Generics to hold two objects of arbitrary types, similar to a *tuple* type found in other languages.

Listing 8.5 The `Pair` Class

```
1 package org.osgi.book.utils;
2
3 public class Pair<A, B> {
4     private final A first;
5     private final B second;
6
7     public Pair(A first, B second) {
8         this.first = first;
9         this.second = second;
10    }
11
12    public A getFirst() {
13        return first;
14    }
15
16    public B getSecond() {
17        return second;
18    }
19
20    // Omitted: hashCode, equals and toString implementations
21 }
```

8.2 Inspecting Headers

The previous example is already very convenient for help providers, but a little inflexible. What if the `help` subdirectory of a bundle is already used for something else, or if we want to give the index file a name other than `index.properties`? It would be helpful if we could somehow provide the path to the index file, so we didn't have to assume a fixed path. But where can we put this information? If we put it in another properties file, then we still need to find *that* file.

However, there is already one file that must appear in the bundle at a fixed location, and that is `META-INF/MANIFEST.MF`. It is also flexible: we are free to add new headers to the manifest so long as their name does not clash with an existing Java or OSGi header. So, we can ask help providers to insert a reference to their index file using a new header name that we define. The provider's manifest might therefore look like something like Listing 8.6.

Fortunately we don't need to parse the manifest ourselves because just like the contents of a bundle, manifest headers are accessible through the `Bundle` interface. In this case we can call the `getHeaders` method to get a dictionary

Listing 8.6 MANIFEST.MF for a Help Provider

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-SymbolicName: org.foo.help
Bundle-Version: 1.0.0
Help-Index: docs/help.properties
```

of all the headers. We can then use this to pick out the value of the `Help-Index` header — note that all the other headers can also be read using this method, including those defined by OSGi and others.

Listing 8.7 shows yet another version of the scanning method. There is still one assumption implicit in this code: the HTML documents are expected to be in the same directory as the index file. However we could remove this assumption by adding more information to either the index or the value of the `Help-Index` header. It might also help to use a more structured file format for the index, such as XML or JSON.

8.3 Tracking Bundles

Somewhere we need to actually call one of the above scanner methods, and do something with the results. Since they all take a `Bundle` object as input, it seems we need to iterate over the set of current bundles. But we should know by now that this is not enough: we also need to register a listener so we can be notified of bundles that are added or removed later. Fortunately this is simple.

Before jumping into the code, let's consider again the lifecycle of a bundle, which we first saw in Section 2.8 and Figure 2.4. Should we allow a bundle in *any* state to contribute help documents, or should we restrict ourselves to a subset of states? This is an important question, for which all extender bundles need to come up with a good answer. Bundle states are not completely exclusive: some bundle states can be considered to “include” other states. For example an `ACTIVE` bundle is also `RESOLVED` and `INSTALLED`. In other words, an extender that is interested in `RESOLVED` bundles should be interested in `STARTING`, `ACTIVE` and `STOPPING` bundles as well. The inclusion relationships of all the states are illustrated in Figure 8.1.

If we include bundles that are merely `INSTALLED` then *every* bundle will be involved in our scan, so users needn't do anything at all beyond installing a bundle in order for it to appear in the index. But this can have negative consequences. For example, some bundles may not be able to resolve because of missing dependencies or other unsatisfied constraints, but they will still be picked by an extender that includes the `INSTALLED` state. This isn't always

Listing 8.7 Scanning a Bundle for Help Documents with Titles (2)

```

98     private static final String HELP_INDEX_BUNDLE_HEADER = "Help-Index";

100    private List<Pair<URL, String>> scanForHelpDocsWithTitle(
101        Bundle bundle) throws IOException, HelpScannerException {
102        @SuppressWarnings("unchecked")
103        Dictionary<String, String> headers = bundle.getHeaders();

105        // Find the index file entry; exit if not found
106        String indexPath = headers.get(HELP_INDEX_BUNDLE_HEADER);
107        if (indexPath == null) {
108            return Collections.emptyList();
109        }
110        URL indexEntry = bundle.getEntry(indexPath);
111        if (indexEntry == null) {
112            throw new HelpScannerException("Entry not found: "
113                + indexPath);
114        }

116        // Calculate the directory prefix
117        int slashIndex = indexPath.lastIndexOf('/');
118        String prefix = (slashIndex == -1)
119            ? "" : indexPath.substring(0, slashIndex);

121        // Load the index file as a Properties object
122        Properties indexProps = new Properties();
123        InputStream stream = null;
124        try {
125            stream = indexEntry.openStream();
126            indexProps.load(stream);
127        } finally {
128            if (stream != null)
129                stream.close();
130        }

132        // Iterate through the files
133        List<Pair<URL, String>> result =
134            new ArrayList<Pair<URL, String>>(indexProps.size());
135        Enumeration<?> names = indexProps.propertyNames();
136        while (names.hasMoreElements()) {
137            String name = (String) names.nextElement();
138            String title = indexProps.getProperty(name);

140            URL entry = bundle.getEntry(prefix + "/" + name + ".html");
141            if (entry != null) {
142                result.add(new Pair<URL, String>(entry, title));
143            }
144        }
145        return result;
146    }

```

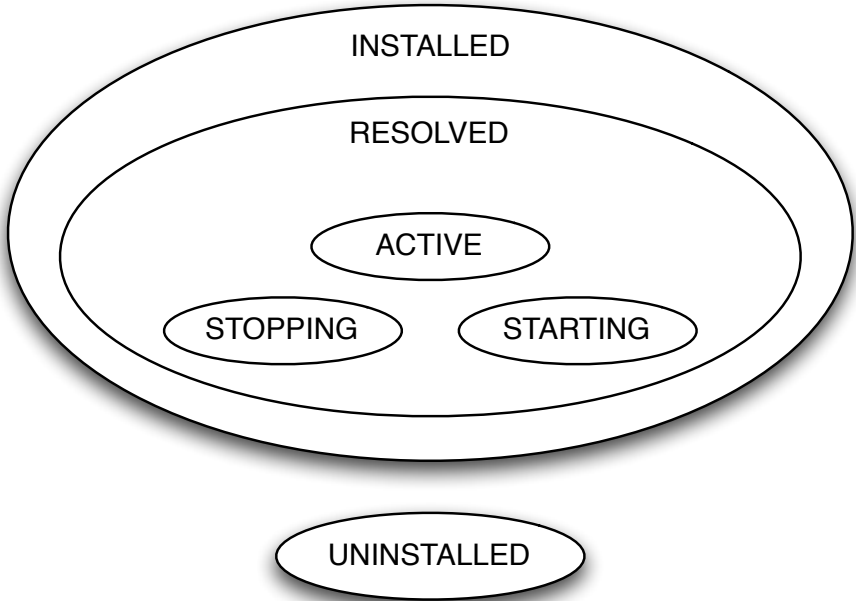


Figure 8.1: Inclusion Relationships of Bundle States

a problem, but in our help system it might be. Merely providing documents doesn't need any dependencies, but suppose our documentation is embedded in a bundle that provides actual functionality as well. If the functionality of the bundle is not available due to a missing dependency, then it could be confusing to display help for it.

So, as an alternative, we can exclude the `INSTALLED` state and include only `RESOLVED` (along with `STARTING`, `ACTIVE` and `STOPPING`) bundles. This ensures only bundles with satisfied dependencies are included by our extender.

But both `INSTALLED` and `RESOLVED` have another problem: it is difficult to exit from those states. Suppose we want the ability for certain bundles to be explicitly removed from consideration by an extender. If the extender is looking for `INSTALLED` bundles, there is obviously only one way to do this: uninstall the bundle. But if the extender is looking for `RESOLVED` bundles, we are still required to uninstall the bundle. The reason for this is there is no explicit “unresolve” method or command in OSGi. We cannot ask for the bundle to be moved from `RESOLVED` back to `INSTALLED`, so we have to use the heavy-handed approach of uninstalling it completely.

The last alternative is to look at `ACTIVE` bundles only. This has the advantage that we can easily move bundles in and out of consideration by the extender simply by starting and stopping them. For this reason most examples of the extender pattern use `ACTIVE` state (with one notable exception, as we will see in Section 8.5).

Listing 8.8 shows the full code for the help system extender (with the exception of the scanning method which was already given in Listing 8.7). It uses `ACTIVE` as the selected bundle state. This class must be constructed with a `BundleContext`, since it needs access to OSGi APIs, and it has three methods in its public API:

- `open` causes the extender to start listening to bundle events, and also scans the existing installed bundles.
- `close` causes the extender to stop listening to bundle events.
- `listHelpDocs` returns a snapshot of all the currently available help documents, as a list of URLs and titles.

Notice how we restrict the extender to including only `ACTIVE` bundles. In the `open` method, we explicitly filter out all bundles that are not in the `ACTIVE` state when we iterate through them. But if we are writing an extender that considers `RESOLVED` state, we would have to select bundles not just in `RESOLVED` state but also in `STARTING`, `ACTIVE` and `STOPPING`, since those states are considered to “include” the `RESOLVED` state. There is a neat trick for testing against multiple states at once: because all of the integer con-

Listing 8.8 The HelpExtender Class

```
1 package org.osgi.book.help.extender;
2
3 import java.io.IOException;
4 import java.io.InputStream;
5 import java.net.URL;
6 import java.util.*;
7 import java.util.concurrent.ConcurrentHashMap;
8
9 import org.osgi.book.utils.Pair;
10 import org.osgi.framework.*;
11 import org.osgi.service.log.LogService;
12 import org.osgi.util.tracker.ServiceTracker;
13
14 public class HelpExtender implements SynchronousBundleListener {
15
16     private final Map<Long, List<Pair<URL, String>>> documentMap
17         = new ConcurrentHashMap<Long, List<Pair<URL, String>>>();
18
19     private final BundleContext context;
20     private final ServiceTracker logTracker;
21
22     public HelpExtender(BundleContext context) {
23         this.context = context;
24         logTracker = new ServiceTracker(context,
25             LogService.class.getName(), null);
26     }
27
28     public void open() {
29         logTracker.open();
30         context.addBundleListener(this);
31
32         Bundle[] bundles = context.getBundles();
33         for (Bundle bundle : bundles) {
34             if (Bundle.ACTIVE == bundle.getState()) {
35                 addingBundle(bundle);
36             }
37         }
38     }
39
40     public List<Pair<URL, String>> listHelpDocs() {
41         List<Pair<URL, String>> result =
42             new ArrayList<Pair<URL, String>>();
43         for (List<Pair<URL, String>> list : documentMap.values()) {
44             result.addAll(list);
45         }
46         return result;
47     }
48
49     public void close() {
50         context.removeBundleListener(this);
51         logTracker.close();
52     }
53
54     /**
55      * Implements BundleListener
56      */
57     public void bundleChanged(BundleEvent event) {
58         if (BundleEvent.STARTED == event.getType()) {
59             addingBundle(event.getBundle());
60         } else if (BundleEvent.STOPPED == event.getType()) {
61             removedBundle(event.getBundle());
62         }
63     }
64 }
```

Listing 8.8 (continued)

```
65 private void addingBundle(Bundle bundle) {
66     long id = bundle.getBundleId();

68     try {
69         List<Pair<URL, String>> docs =
70             scanForHelpDocsWithTitle(bundle);
71         if (!docs.isEmpty()) {
72             documentMap.put(id, docs);
73         }
74     } catch (IOException e) {
75         logError("IO error in bundle " + bundle.getLocation(), e);
76     } catch (HelpScannerException e) {
77         logError("Error in bundle " + bundle.getLocation(), e);
78     }
79 }

81 private void removedBundle(Bundle bundle) {
82     documentMap.remove(bundle.getBundleId());
83 }

85 private void logError(String message, Throwable t) {
86     LogService log = (LogService) logTracker.getService();
87     if (log != null) {
88         log.log(LogService.LOG_ERROR, message, t);
89     } else {
90         System.err.println(message);
91         if (t != null) t.printStackTrace();
92     }
93 }

95 // Omitted: scanForHelpDocsWithTitle method from previous section
96 // ...
```

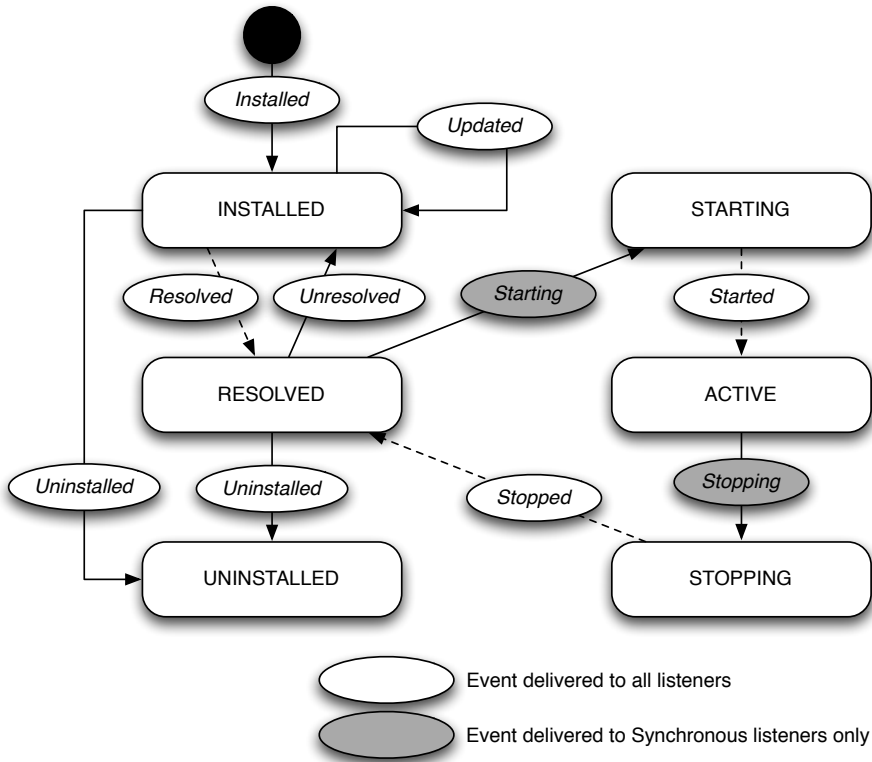


Figure 8.2: Bundle Transitions and Events

stants representing states are powers of two, we can use a bitwise calculation as shown in Listing 8.9

Listing 8.9 Testing Multiple Bundle States in One if Statement

```
if (bundle.getState() & (RESOLVED | STARTING | ACTIVE | STOPPING) > 0) {  
    addingBundle(bundle);  
}
```

The other way we ensure only active bundles are considered is by looking at the type of the bundle event received by the `bundleChanged` method. This indicated which transition the bundle state has just taken, so we match the two transitions which take us into (`STARTED`) and out of (`STOPPED`) the state we are interested in.

Figure 8.2 shows the events that are fired on each transition. One peculiarity is that when an `RESOLVED` bundle is uninstalled, it does not fire an UN-

RESOLVED event, but instead goes straight to UNINSTALLED, firing only an UNINSTALLED event. So extenders that take an interest in RESOLVED bundles must watch for both UNRESOLVED and UNINSTALLED events in order to know when to drop a bundle — bearing in mind that some of the UNINSTALLED events should be ignored, since they come from bundle that were never in the RESOLVED state at all.

Note that the private method `addingBundle` is called both by the initial scan of bundles during `open` and also by the event callback method, `bundleChanged`. This method must be able to handle duplicate calls with the same bundle, since if a bundle is activated while `open` is running, between the call to `addBundleListener` and `getBundles`, that bundle will be processed twice. In this case our use of a map means the second update simply overwrites the first, which is fine.

To test this extender, we will write a shell command that prints a list of the available help documents; this is shown in Listing 8.10. The corresponding activator is in Listing 8.11 and the `bnf` descriptor is in Listing 8.12.

If we install and start the resulting bundle, we should now be able to call the `helpDocs` command:

```
-> helpDocs
0 documents(s) found
```

Of course, we haven't built a bundle yet that provides any help documentation. Lets do that now. If we create a directory in our project called `resources/help_sample`, we can put our index file and help HTML there and include them in a bundle using the `Include-Resource` instruction to `bnf`, as shown in Listing 8.13. Note the assignment-style format of the `Include-Resource` header; this tells `bnf` to copy the contents of the directory on the right hand side into a directory named `docs` in the bundle. If we omitted `"docs="` then the files would be copied to the root of the bundle.

Now if we install and start the `help_sample` bundle and rerun the `helpDocs` command, we should see the following output:

```
-> helpDocs
4 document(s) found
Introduction (bundle://12.0:0/docs/introduction.html)
First Steps in OSGi (bundle://12.0:0/docs/first_steps.html)
Bundle Dependencies (bundle://12.0:0/docs/dependencies.html)
Introduction to Services (bundle://12.0:0/docs/intro_services.html)
```

8.4 Synchronous and Asynchronous Bundle Listeners

In the previous section we wrote our extender as an implementation of `BundleListener`. But there is also another interface for listening to bundle events: `Syn-`

Listing 8.10 Shell Command for Testing the Help Extender

```
1 package org.osgi.book.help.extender;
2
3 import java.io.PrintStream;
4 import java.net.URL;
5 import java.util.List;
6
7 import org.apache.felix.shell.Command;
8 import org.osgi.book.utils.Pair;
9
10 public class HelpListCommand implements Command {
11
12     private final HelpExtender extender;
13
14     public HelpListCommand(HelpExtender extender) {
15         this.extender = extender;
16     }
17
18     public void execute(String line, PrintStream out,
19         PrintStream err) {
20         List<Pair<URL, String>> docs = extender.listHelpDocs();
21         out.println(docs.size() + " document(s) found");
22         for (Pair<URL, String> pair : docs) {
23             out.println(pair.getSecond() + " (" + pair.getFirst() + ")");
24         }
25     }
26
27     public String getName() {
28         return "helpDocs";
29     }
30
31     public String getShortDescription() {
32         return "List currently available help documentation.";
33     }
34
35     public String getUsage() {
36         return "helpDocs";
37     }
38
39 }
```

Listing 8.11 Activator for the Help Extender Bundle

```

1 package org.osgi.book.help.extender;

3 import org.apache.felix.shell.Command;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6 import org.osgi.framework.ServiceRegistration;

8 public class HelpExtenderActivator implements BundleActivator {

10     private volatile HelpExtender extender;
11     private volatile ServiceRegistration cmdSvcReg;

13     public void start(BundleContext context) throws Exception {
14         extender = new HelpExtender(context);
15         extender.open();

17         HelpListCommand command = new HelpListCommand(extender);
18         cmdSvcReg = context.registerService(Command.class.getName(),
19             command, null);
20     }

22     public void stop(BundleContext context) throws Exception {
23         cmdSvcReg.unregister();
24         extender.close();
25     }

27 }

```

Listing 8.12 Bnd Descriptor for the Help Extender Bundle

```

# helpextender.bnd
Private-Package: org.osgi.book.help.extender,\
    org.osgi.book.utils
Bundle-Activator: org.osgi.book.help.extender.HelpExtenderActivator

```

Listing 8.13 Bnd Descriptor for a Sample Help Provider

```

# help_sample.bnd
Help-Index: docs/index.properties
Include-Resource: docs=resources/help_sample
Import-Package:

```

chronousBundleListener.

The `SynchronousBundleListener` interface (“SBL”) is actually a sub-type of `BundleListener` (“BL”) and adds no methods of its own, making it a marker interface. So the real difference between the two is how the OSGi framework delivers events to them.

SBLs are the simplest to understand: when the state of a bundle changes, the framework will call methods on all registered SBLs synchronously (i.e. in the same thread) before returning to the original caller. This is illustrated in Figure 8.3, which shows a UML-like sequence diagram for starting a bundle in the presence of SBLs. First the listeners are notified that the bundle is STARTING; next the bundle is actually started by calling `start` on its activator; and finally the listeners are notified that the bundle has STARTED.

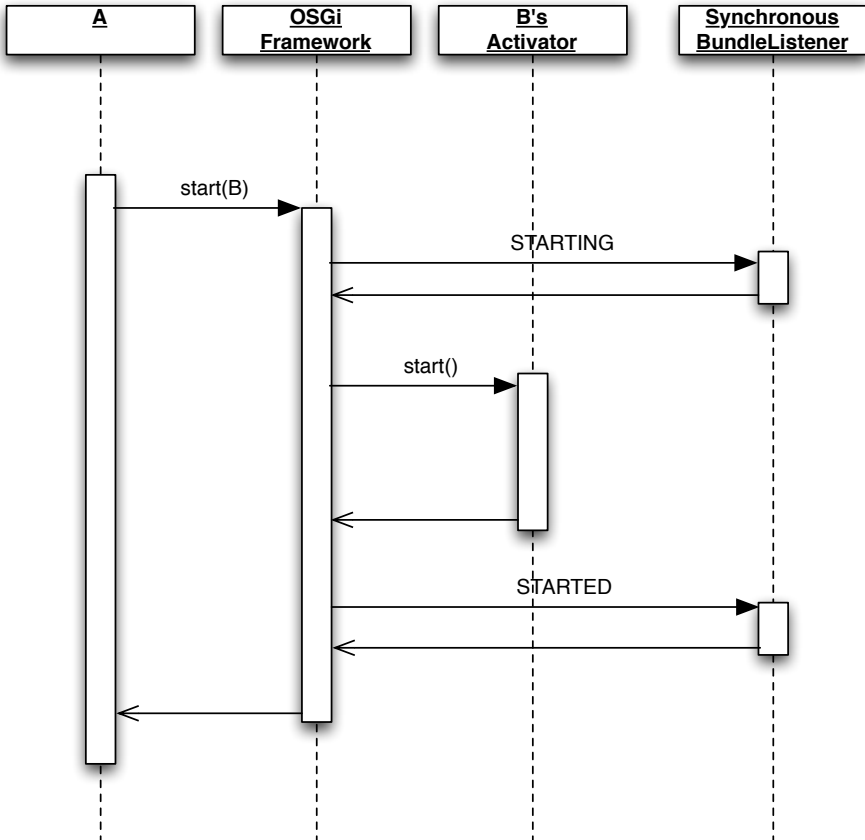


Figure 8.3: Synchronous Event Delivery when Starting a Bundle

BLs, in contrast, do not receive bundle events immediately. Instead, events are




Figure 8.4: Asynchronous Event Delivery after Starting a Bundle

added to a queue, and delivered a little while later by one or more background threads responsible for processing the queue. Unfortunately it is impossible to know exactly how much later the events will be received. Also some of the transient events are simply not delivered, such as `STARTING` and `STOPPING`. This is illustrated in Figure 8.4.

The main advantage of BLs is they are easier to program. The framework guarantees never to call the `bundleChanged` method concurrently from multiple threads, so we do not need to worry about locking fields for atomicity — however it does *not* guarantee to always call `bundleChanged` from the *same* thread, so we still need to follow safe publication idioms as described in Section 6.3. Also, since the callback to a BL is executing in a thread dedicated to that purpose rather than directly in an arbitrary caller’s thread, we can be a little more liberal about performing blocking operations and other computations that would be too long-running for execution in an SBL. We still need to be cautious though, since we could hold up the delivery of bundle events to other BLs, so truly long-running operations should be performed in an thread that we explicitly create.

On the other hand, because delivery of events to a BL is delayed, they can be stale. For example, if we had built our Help extender as a BL rather than an SBL, we would only hear about a bundle being uninstalled sometime after

it was uninstalled. Therefore a user might request a help document from a bundle which no longer exists. This would lead to an error which we would need to catch and display appropriately to the user.

8.5 The Eclipse Extension Registry

As mentioned in Section ??, the Eclipse IDE and platform are based on OSGi. However, Eclipse currently makes very little use of services, mainly for historical reasons.

Eclipse did not always use OSGi: in fact it only started using OSGi in version 3.0, which was released in 2004, nearly three years after the first release. Until version 3.0, Eclipse used its own custom-built module system which was somewhat similar to OSGi, but substantially less powerful and robust. Also it used its own late-binding mechanism called the “extension registry”, which achieves roughly the same goal as OSGi’s services but in a very different way. When Eclipse switched to OSGi, it threw out the old module system, but it did *not* throw out the extension registry, because to do so would have rendered almost all existing Eclipse plug-ins useless. By that time there were already many thousands of plug-ins for Eclipse, and there was no feasible way to offer a compatibility layer that could commute extension registry based code into services code. Therefore the extension registry continues to be the predominant model for late-binding in Eclipse plug-ins and RCP applications.

Today, the extension registry is implemented as an extender bundle. Bundles are able to declare both *extension points* and *extensions* through the special file `plugin.xml` which must appear at the root of a bundle. This file used to be the central file of Eclipse’s old module system, as it listed the dependencies and exports of a plug-in, but today its role is limited to the extension registry¹

Listing 8.14 shows an example of a `plugin.xml` file, which in this case declares both an extension point and an extension in the same bundle. An *extension point* is a place where functionality can be contributed, and an *extension* is the declaration of that contributed functionality. They are like a socket and a plug, respectively. In this example the extension contributes functionality into the `org.foo.commands` extension point, which is defined in the *same* bundle: there is nothing to stop this and it can be useful in certain situations.

Usually one does not directly edit `plugin.xml` as XML in a text editor. Instead there are powerful tools in Eclipse PDE to edit the extensions and extension points graphically. They are shown in Figures 8.5 and 8.6.

We will not go into much detail on how the extension registry is used, as this subject is well documented in various books and articles about Eclipse.

¹One still occasionally sees a `plugin.xml` that includes the old module system declarations; they are supported by a compatibility layer.

Listing 8.14 An Eclipse plugin.xml File

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <?eclipse version="3.2"?>
3 <plugin>

5     <extension-point id="org.foo.commands"
6                     name="Commands"
7                     schema="schema/org.foo.commands.exsd"/>

9     <extension point="org.foo.commands">
10         <command id="org.bar.mycommand"
11                 class="org.bar.MyCommand">
12             </command>
13     </extension>

15 </plugin>

```

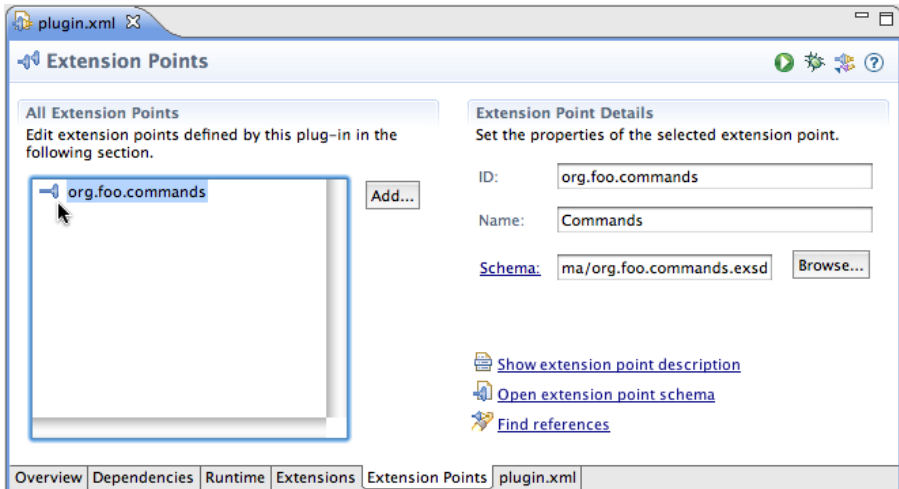


Figure 8.5: Editing an Extension Point in Eclipse PDE

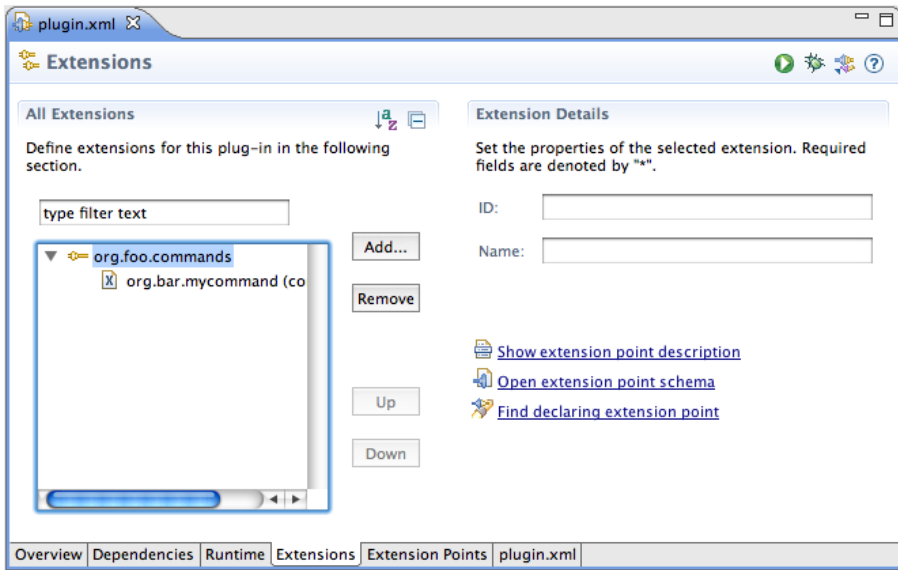


Figure 8.6: Editing an Extension in Eclipse PDE

However we should already be able to guess how the registry bundle works by scanning bundles for the presence of a `plugin.xml` file, reading the extension and extension point declarations therein, and merging them into a global map. It is not really so different from the help extender that we implemented in the previous section.

However, rather than considering only `ACTIVE` bundles, the extension registry considers all `RESOLVED` bundles. The result — as we would expect from the discussion in Section 8.3 — is that we don't have a lot of control over the content of the registry: *all* extensions in all `RESOLVED` bundles will contribute to the running application, and the only way we can remove those contributed features is by fully uninstalling the bundle.

8.6 Impersonating a Bundle

In Section 2.10 we learned that the only way to obtain a `BundleContext` and interact with the framework was by implementing a `BundleActivator`. However we also learned in a footnote that that was not strictly true: another way is to define an extender bundle that interacts with the framework *on behalf of* another bundle. By doing this we can define standard patterns of interaction with the framework.

For example, suppose we have a lot of bundles that register Mailbox services.

The activator code to create the Mailbox instances and register them with the service registry can be quite repetitive, and it might be useful to use the extender model to define an alternative based on declarations. Thus we could offer a Mailbox service simply by adding to our `MANIFEST.MF`:

```
Mailbox-ServiceClass: org.example.MyMailbox
```

But there is a catch: if we register the Mailbox services using the `BundleContext` of our extender bundle, then they will all appear to be services offered by that extender, not by the real bundle that contains the Mailbox implementation. This would be wrong. However, there is a solution: in OSGi Release 4.1 a new method was added to the `Bundle` interface called `getBundleContext`. This method allows our extender to register services as if they were registered by the target bundle — in other words it allows us to impersonate bundles.

It should be noted that this is an advanced technique. There are many factors that need to be taken into account when writing an impersonating extender, so it should *not* be undertaken lightly. Also note that the example in this section is simply a limited version of the standard OSGi facility called Declarative Services (which we will look at in detail Chapter ??) so we would not wish to do this for real. Nevertheless, it is useful to understand what is going on when *using* an impersonating extender.

To ease the implementation of this extender, we will use a helper class called `BundleTracker`. One problem we have seen in our extender code so far is that we must start listening to bundle events before getting the list of current bundles. This is necessary to avoid missing any events, but it means we can get duplicates. So in each example so far we have written our code cautiously so that it works correctly even if called twice with the same bundle. `BundleTracker` simplifies the code for working with bundle events in the same way that `ServiceTracker` simplifies working with service events: it guarantees to call `addingBundle` and `removedBundle` exactly once for each bundle. Unfortunately there is no `BundleTracker` in the OSGi core or compendium APIs², so we will use our own. The code for this is given in Appendix A.

Listing 8.15 shows the code for the extender. Although it is quite long, the bulk is taken up with using Java reflection to load the specified Mailbox class by name, instantiating it, and handling the myriad checked exceptions that Java gives us whenever we use reflection. Note that we must ask the *target* bundle to perform the class load, by calling the `Bundle.loadClass` method, since in general the extender bundle will not know anything about the Mailbox implementation class. Also we assume that the specified class has a zero-argument constructor; if that is not the case then we will get an `InstantiationException`.

²This is true as of Release 4.1, however future releases of OSGi may include a standard `BundleTracker`.

Listing 8.15 Mailbox Service Extender

```

1 package org.osgi.book.extender.service;
2
3 import java.util.HashMap;
4 import java.util.Map;
5
6 import org.osgi.book.reader.api.Mailbox;
7 import org.osgi.framework.Bundle;
8 import org.osgi.framework.BundleContext;
9 import org.osgi.framework.ServiceRegistration;
10 import org.osgi.service.log.LogService;
11 import org.osgi.utils.BundleTracker;
12
13 public class MailboxServiceExtender extends BundleTracker {
14
15     private static final String SVC_HEADER = "Mailbox-ServiceClass";
16
17     private final Map<String, ServiceRegistration> registrations
18         = new HashMap<String, ServiceRegistration>();
19     private final LogService log;
20
21     public MailboxServiceExtender(BundleContext ctx, LogService log) {
22         super(ctx);
23         this.log = log;
24     }
25
26     @Override
27     protected void addingBundle(Bundle bundle) {
28         String className = (String) bundle.getHeaders().get(SVC_HEADER);
29
30         if (className != null) {
31             try {
32                 Class<?> svcClass = bundle.loadClass(className);
33                 if (!Mailbox.class.isAssignableFrom(svcClass)) {
34                     log.log(LogService.LOG_ERROR,
35                         "Declared class is not an instance of Mailbox");
36                     return;
37                 }
38                 Object instance = svcClass.newInstance();
39                 ServiceRegistration reg = bundle.getBundleContext()
40                     .registerService(Mailbox.class.getName(), instance,
41                         null);
42                 synchronized (registrations) {
43                     registrations.put(bundle.getLocation(), reg);
44                 }
45             } catch (ClassNotFoundException e) {
46                 log.log(LogService.LOG_ERROR, "Error creating service", e);
47             } catch (InstantiationException e) {
48                 log.log(LogService.LOG_ERROR, "Error creating service", e);
49             } catch (IllegalAccessException e) {
50                 log.log(LogService.LOG_ERROR, "Error creating service", e);
51             }
52         }
53     }
54
55     @Override
56     protected void removedBundle(Bundle bundle) {
57         ServiceRegistration reg;
58         synchronized (registrations) {
59             reg = registrations.remove(bundle.getLocation());
60         }
61         if (reg != null) {
62             reg.unregister();
63         }
64     }
65 }

```

The activator for this extender is shown in Listing 8.16 and the `bnd` descriptor is in Listing 8.17. In order to give the extender itself access to an instance of `LogService`, we use the `LogTracker` class from Section 4.10.1.

Listing 8.16 Activator for the Mailbox Service Extender

```

1 package org.osgi.book.extender.service;
2
3 import org.osgi.book.utils.LogTracker;
4 import org.osgi.framework.BundleActivator;
5 import org.osgi.framework.BundleContext;
6
7 public class MailboxServiceExtenderActivator
8     implements BundleActivator {
9
10    private volatile LogTracker logTracker;
11    private volatile MailboxServiceExtender extender;
12
13    public void start(BundleContext context) throws Exception {
14        logTracker = new LogTracker(context);
15        logTracker.open();
16
17        extender = new MailboxServiceExtender(context, logTracker);
18        extender.open();
19    }
20
21    public void stop(BundleContext context) throws Exception {
22        extender.close();
23        logTracker.close();
24    }
25 }
26 }

```

Listing 8.17 Bnd Descriptor for the Mailbox Service Extender

```

# mbox_svc_extender.bnd
Private-Package: org.osgi.book.utils,\
    org.osgi.book.tracker,\
    org.osgi.book.extender.service
Bundle-Activator:\
    org.osgi.book.extender.service.MailboxServiceExtenderActivator

```

Now we can write a simple bundle that contains a minimal Mailbox implementation and declares it through a declaration in its `MANIFEST.MF`. Listing 8.18 shows the mailbox implementation we will use and Listing 8.19 is the `bnd` descriptor.

When we install and start the `sample_svc_extender` bundle, we should be able to see the registered mailbox service by typing the `services` command:

```

-> services
...
sample_svc_extender (6) provides:
-----
org.osgi.book.reader.api.Mailbox

```

Listing 8.18 Minimal Mailbox Class

```
1 package org.osgi.book.extender.service.sample;
2
3 import org.osgi.book.reader.api.*;
4
5 public class MyMailbox implements Mailbox {
6
7     public long [] getAllMessages () {
8         return new long [0];
9     }
10    public Message [] getMessages(long [] ids) {
11        return new Message [0];
12    }
13    public long [] getMessagesSince(long id) {
14        return new long [0];
15    }
16    public void markRead(boolean read, long [] ids) {
17    }
18 }
```

Listing 8.19 Bnd Descriptor for a “Declarative” Mailbox Service Bundle

```
# sample_svc_extender.bnd
Private-Package: org.osgi.book.extender.service.sample
Mailbox-ServiceClass: org.osgi.book.extender.service.sample.MyMailbox
```

Here we see that a Mailbox service has been registered, apparently by the `sample_svc_extender` bundle, although it was really registered by the extender on behalf of that bundle.

To reiterate, impersonating other bundles is an advanced technique that must be used with care. In some cases it can be used to implement common patterns in a single location; however the example in this chapter could also have been implemented with straightforward Java inheritance, i.e. a simple activator that can be used as a base class.

8.7 Conclusion

The extender model is a very useful technique for allowing an application to be extended via resources or declarations, rather than via programmatic services. In some cases extenders can be used to implement common bundle implementation patterns by acting on behalf of other bundles.

9 Configuration and Metadata

TODO

Part II

**Component Oriented
Development**

10 Component Oriented Development

TODO

11 Declarative Services

TODO

Part III

Practical OSGi

12 Testing OSGi Bundles

TODO

13 Using Third-Party Libraries

TODO

14 Building Web Applications

TODO

Part IV

Appendices

A Bundle Tracker

The following is an example implementation of a Bundle Tracker. Note that Release 4.2 of the OSGi specification will include a standardised implementation of `BundleTracker` with a slightly different (though very similar) API.

This class tracks only bundles in ACTIVE state.

Listing A.1 BundleTracker

```
1 package org.osgi.utils;
2
3 import java.util.Collection;
4 import java.util.HashMap;
5 import java.util.Map;
6
7 import org.osgi.framework.Bundle;
8 import org.osgi.framework.BundleContext;
9 import org.osgi.framework.BundleEvent;
10 import org.osgi.framework.BundleListener;
11 import org.osgi.framework.SynchronousBundleListener;
12
13 public abstract class BundleTracker {
14
15     private final BundleContext context;
16     private final Map<String, Bundle> trackedSet
17         = new HashMap<String, Bundle>();
18
19     private final BundleListener listener =
20         new SynchronousBundleListener() {
21         public void bundleChanged(BundleEvent event) {
22             if (BundleEvent.STARTED == event.getType()) {
23                 internalAdd(event.getBundle());
24             } else if (BundleEvent.STOPPING == event.getType()) {
25                 internalRemove(event.getBundle());
26             }
27         }
28     };
29
30     private boolean isOpen = false;
31
32     public BundleTracker(BundleContext context) {
33         this.context = context;
34     }
35
36     /**
37      * Open this BundleTracker and begin tracking
38      */
39     public final void open() {
40         synchronized (this) {
41             if (isOpen) {
42                 return;
43             }
44         }
45     }
46 }
```

Listing A.1 (continued)

```

43         } else {
44             isOpen = true;
45         }
46     }

48     context.addBundleListener(listener);

50     Bundle[] bundles = context.getBundles();
51     for (Bundle bundle : bundles) {
52         if (Bundle.ACTIVE == bundle.getState()) {
53             internalAdd(bundle);
54         }
55     }
56 }

58 /**
59  * Close this <code>BundleTracker</code> and stop tracking
60  */
61 public final void close() {
62     synchronized (this) {
63         if (!isOpen) {
64             return;
65         } else {
66             isOpen = false;
67         }
68     }
69     context.removeBundleListener(listener);

71     Bundle[] bundles;
72     synchronized (trackedSet) {
73         bundles = trackedSet.values().toArray(
74             new Bundle[trackedSet.size()]);
75     }

77     for (Bundle bundle : bundles) {
78         internalRemove(bundle);
79     }
80 }

82 /**
83  * Return a snapshot of the currently tracked bundles.
84  */
85 public Bundle[] getBundles() {
86     synchronized (trackedSet) {
87         Collection<Bundle> bundles = trackedSet.values();
88         Bundle[] result = bundles.toArray(
89             new Bundle[bundles.size()]);
90     }
91     return result;
92 }

94 /**
95  * Called when a bundle is being added to the
96  * <code>BundleTracker</code>. This method does nothing, it is
97  * expected to be overridden by concrete subclasses.
98  *
99  * @param bundle
100  */
101 protected void addingBundle(Bundle bundle) {
102 }

```

Listing A.1 (continued)

```
104  /**
105   * Called when a bundle is being removed from the
106   * BundleTracker This method does nothing, it is
107   * expected to be overridden by concrete subclasses.
108   *
109   * @param bundle
110   */
111  protected void removedBundle(Bundle bundle) {
112  }

114  private void internalAdd(Bundle bundle) {
115      Bundle prior;
116      synchronized (trackedSet) {
117          prior = trackedSet.put(bundle.getLocation(), bundle);
118      }
119      if (prior == null) {
120          addingBundle(bundle);
121      }
122  }

124  private void internalRemove(Bundle bundle) {
125      Bundle removed;
126      synchronized (trackedSet) {
127          removed = trackedSet.remove(bundle.getLocation());
128      }
129      if (removed != null) {
130          removedBundle(removed);
131      }
132  }
133 }
```

B ANT Build System for Bnd

The following is suggested project structure for building OSGi projects based on ANT and bnd. The project structure is assumed to be as in Figure B.1. The purpose of these directories is as follows:

src contains the Java source of our bundles, laid out in the normal Java way with subdirectories for each package.

test contains JUnit-based tests for our Java source, also laid out in package subdirectories.

bundles contains binary or pre-built bundles as JARs that form the dependencies of our code. We will need these at compile time as well as runtime. For example we may include `osgi.cmpn.jar`, which is a bundle that contains the API (but *not* implementation!) of all the OSGi Compendium services.

All of the bnd descriptor files are placed at the top level of the project¹. This is also where we place `build.xml` and a supplementary properties file called `build.properties`. The latter file is shown in Listing B.1; the settings shown will certainly need to be changed to match your own computer.

Listing B.1 `build.properties`

```
1 # Path to the Felix installation directory
2 felix.home=/path/to/felix-1.0.3
3
4 # Location of the JUnit JAR
5 junit.path=/path/to/junit/junit-4.4.jar
6
7 #Location of bnd.jar
8 bnd.path=/path/to/bnd/bnd.jar
```

¹This could be changed to a subdirectory by editing the `bundle` target of the ANT build.

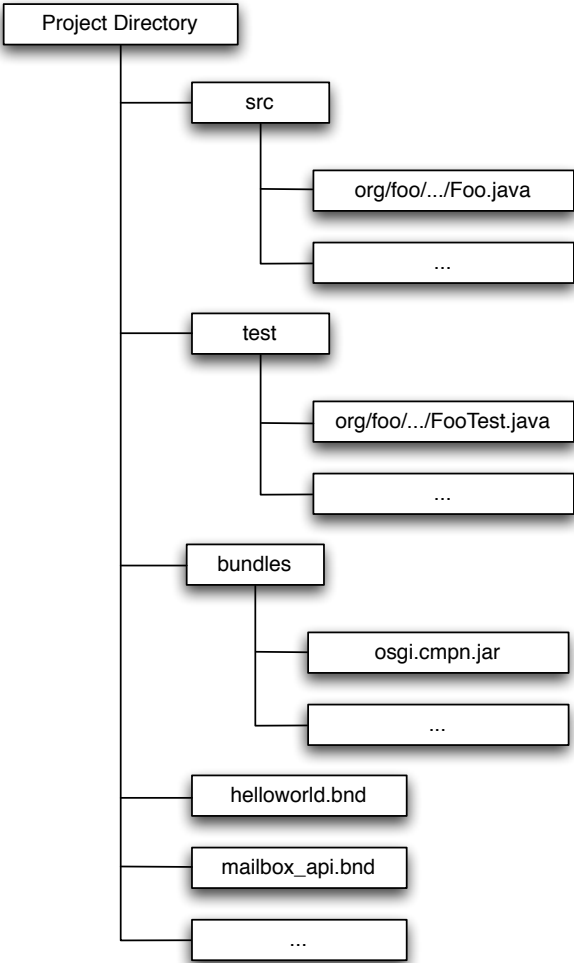


Figure B.1: OSGi Project Structure

Listing B.2 build.xml

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <project name="osgibook" default="bundle">

4     <!-- Import machine-specific settings -->
5     <property file="build.properties"/>

7     <!-- Setup build paths -->
8     <property name="build_dir" value="build"/>
9     <property name="build_classes_dir" value="${build_dir}/classes"/>
10    <property name="build_bundles_dir" value="${build_dir}/bundles"/>
11    <property name="build_test_dir" value="${build_dir}/tests"/>

13    <!-- Set a classpath for the OSGi libraries -->
14    <path id="osgilibs">
15        <pathelement location="${felix.home}/bin/felix.jar"/>
16        <fileset dir="bundles" includes="*.jar"/>
17    </path>

19    <!-- Set a classpath for JUnit tests -->
20    <path id="test_classpath">
21        <path refid="osgilibs"/>
22        <pathelement location="${junit.path}"/>
23    </path>

25    <!-- Load the bnd custom task -->
26    <taskdef resource="aQute/bnd/ant/taskdef.properties"
27        classpath="${bnd.path}"/>

29    <!-- TARGET: clean; cleans all build outputs -->
30    <target name="clean" description="Clean all build outputs">
31        <delete dir="${build_dir}"/>
32    </target>

34    <!-- TARGET: compile; compiles Java sources -->
35    <target name="compile" description="Compile Java sources">
36        <mkdir dir="${build_classes_dir}"/>
37        <javac srcdir="src" destdir="${build_classes_dir}"
38            debug="true" classpathref="osgilibs"/>

40        <mkdir dir="${build_test_dir}"/>
41        <javac srcdir="test" destdir="${build_test_dir}"
42            debug="true" classpathref="test_classpath"/>
43    </target>

45    <!-- TARGET: bundle; generates bundle JARs using bnd -->
46    <target name="bundle" depends="compile"
47        description="Build bundles">
48        <mkdir dir="${build_bundles_dir}"/>
49        <!-- Convert an ANT fileset to a flat list of files -->
50        <pathconvert property="bnd.files" pathsep=" "
51            <fileset dir="${basedir}"
52                <include name="*.bnd"/>
53            </fileset>
54        </pathconvert>
55        <bnd classpath="${build_classes_dir}" failok="false"
56            output="${build_bundles_dir}" files="${bnd.files}"/>
57    </target>
58 </project>

```

Bibliography

- [1] Boeing 747 Fun Facts. http://www.boeing.com/commercial/747family/pf/pf_facts.html.
- [2] Java JAR File Specification. <http://java.sun.com/j2se/1.4.2/docs/guide/jar/jar.html>.
- [3] Apache Jakarta Commons HttpClient. <http://jakarta.apache.org/httpcomponents/httpclient-3.x/>.
- [4] Eclipse Equinox. <http://www.eclipse.org/equinox>.
- [5] Eclipse. Eclipse Public License 1.0. <http://opensource.org/licenses/eclipse-1.0.php>.
- [6] Knopflerfish OSGi. <http://www.knopflerfish.org/>.
- [7] Apache Felix. <http://felix.apache.org/site/index.html>.
- [8] Concierge OSGi. <http://concierge.sourceforge.net/>.
- [9] JSR 277: Java Module System. <http://www.jcp.org/en/jsr/detail?id=277>.
- [10] bnd - Bundle Tool. <http://www.aqute.biz/Code/Bnd>.
- [11] Internet RFC 1738. <http://www.rfc.net/rfc1738.html>.
- [12] Twitter. <http://twitter.com/home>.
- [13] Jaiku. <http://www.jaiku.com/>.
- [14] WinZip® - The Zip File Utility for Windows. <http://www.winzip.com/index.htm>.
- [15] Brian Goetz. *Java Concurrency in Practice*. Addison-Wesley, 2005.
- [16] “Deadlock” on Wikipedia. <http://en.wikipedia.org/wiki/Deadlock>.
- [17] B.A. Botkin and A.F. Harlow. *A Treasury of Railroad Folklore*. Crown Publishers, New York, NY, USA, 1953.
- [18] Edsger W. Dijkstra. Hierarchical ordering of sequential processes. In *Operating Systems Techniques*, pages 72–93. n.d.
- [19] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Boston, MA, USA, 1995.