

Spring Boot Reference Guide

2.0.0.M3

Phillip Webb , Dave Syer , Josh Long , Stéphane Nicoll , Rob Winch , Andy Wilkinson ,
Marcel Overdijk , Christian Dupuis , Sébastien Deleuze , Michael Simons , Vedran Pavi#

Copyright © 2012-2017

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this Copyright Notice, whether distributed in print or electronically.

Table of Contents

I. Spring Boot Documentation	1
1. About the documentation	2
2. Getting help	3
3. First steps	4
4. Working with Spring Boot	5
5. Learning about Spring Boot features	6
6. Moving to production	7
7. Advanced topics	8
II. Getting started	9
8. Introducing Spring Boot	10
9. System Requirements	11
9.1. Servlet containers	11
10. Installing Spring Boot	12
10.1. Installation instructions for the Java developer	12
Maven installation	12
Gradle installation	13
10.2. Installing the Spring Boot CLI	14
Manual installation	14
Installation with SDKMAN!	14
OSX Homebrew installation	15
MacPorts installation	15
Command-line completion	15
Quick start Spring CLI example	16
10.3. Upgrading from an earlier version of Spring Boot	16
11. Developing your first Spring Boot application	17
11.1. Creating the POM	17
11.2. Adding classpath dependencies	18
11.3. Writing the code	19
The @RestController and @RequestMapping annotations	19
The @EnableAutoConfiguration annotation	19
The “main” method	20
11.4. Running the example	20
11.5. Creating an executable jar	20
12. What to read next	22
III. Using Spring Boot	23
13. Build systems	24
13.1. Dependency management	24
13.2. Maven	24
Inheriting the starter parent	24
Using Spring Boot without the parent POM	25
Using the Spring Boot Maven plugin	26
13.3. Gradle	26
13.4. Ant	26
13.5. Starters	27
14. Structuring your code	33
14.1. Using the “default” package	33
14.2. Locating the main application class	33

15. Configuration classes	35
15.1. Importing additional configuration classes	35
15.2. Importing XML configuration	35
16. Auto-configuration	36
16.1. Gradually replacing auto-configuration	36
16.2. Disabling specific auto-configuration	36
17. Spring Beans and dependency injection	37
18. Using the @SpringBootApplication annotation	38
19. Running your application	39
19.1. Running from an IDE	39
19.2. Running as a packaged application	39
19.3. Using the Maven plugin	39
19.4. Using the Gradle plugin	40
19.5. Hot swapping	40
20. Developer tools	41
20.1. Property defaults	41
20.2. Automatic restart	42
Excluding resources	43
Watching additional paths	43
Disabling restart	43
Using a trigger file	44
Customizing the restart classloader	44
Known limitations	44
20.3. LiveReload	45
20.4. Global settings	45
20.5. Remote applications	45
Running the remote client application	46
Remote update	47
21. Packaging your application for production	48
22. What to read next	49
IV. Spring Boot features	50
23. SpringApplication	51
23.1. Startup failure	51
23.2. Customizing the Banner	52
23.3. Customizing SpringApplication	53
23.4. Fluent builder API	53
23.5. Application events and listeners	54
23.6. Web environment	54
23.7. Accessing application arguments	55
23.8. Using the ApplicationRunner or CommandLineRunner	55
23.9. Application exit	56
23.10. Admin features	56
24. Externalized Configuration	57
24.1. Configuring random values	58
24.2. Accessing command line properties	58
24.3. Application property files	59
24.4. Profile-specific properties	60
24.5. Placeholders in properties	61
24.6. Using YAML instead of Properties	61
Loading YAML	61

Exposing YAML as properties in the Spring Environment	62
Multi-profile YAML documents	62
YAML shortcomings	63
Merging YAML lists	63
24.7. Type-safe Configuration Properties	64
Third-party configuration	67
Relaxed binding	67
Properties conversion	68
@ConfigurationProperties Validation	69
@ConfigurationProperties vs. @Value	70
25. Profiles	71
25.1. Adding active profiles	71
25.2. Programmatically setting profiles	71
25.3. Profile-specific configuration files	72
26. Logging	73
26.1. Log format	73
26.2. Console output	73
Color-coded output	74
26.3. File output	75
26.4. Log Levels	75
26.5. Custom log configuration	76
26.6. Logback extensions	77
Profile-specific configuration	78
Environment properties	78
27. Developing web applications	79
27.1. The 'Spring Web MVC framework'	79
Spring MVC auto-configuration	79
HttpMessageConverters	80
Custom JSON Serializers and Deserializers	80
MessageCodesResolver	81
Static Content	81
Custom Favicon	83
ConfigurableWebBindingInitializer	83
Template engines	83
Error Handling	83
Custom error pages	84
Mapping error pages outside of Spring MVC	85
Error Handling on WebSphere Application Server	85
Spring HATEOAS	86
CORS support	86
27.2. The 'Spring WebFlux framework'	86
Spring WebFlux auto-configuration	87
HTTP codecs with HttpMessageReaders and HttpMessageWriters	88
Static Content	88
Template engines	89
27.3. JAX-RS and Jersey	89
27.4. Embedded servlet container support	90
Servlets, Filters, and listeners	90
Registering Servlets, Filters, and listeners as Spring beans	90
Servlet Context Initialization	91

Scanning for Servlets, Filters, and listeners	91
The ServletWebServerApplicationContext	91
Customizing embedded servlet containers	91
Programmatic customization	92
Customizing ConfigurableServletWebServerFactory directly	92
JSP limitations	93
28. Security	94
28.1. OAuth2	95
Authorization Server	95
Resource Server	95
28.2. Token Type in User Info	96
28.3. Customizing the User Info RestTemplate	96
Client	97
Single Sign On	98
28.4. Actuator Security	98
29. Working with SQL databases	100
29.1. Configure a DataSource	100
Embedded Database Support	100
Connection to a production database	101
Connection to a JNDI DataSource	102
29.2. Using JdbcTemplate	102
29.3. JPA and 'Spring Data'	103
Entity Classes	103
Spring Data JPA Repositories	104
Creating and dropping JPA databases	105
Open EntityManager in View	105
29.4. Using H2's web console	105
Changing the H2 console's path	106
Securing the H2 console	106
29.5. Using jOOQ	106
Code Generation	106
Using DSLContext	106
jOOQ SQL dialect	107
Customizing jOOQ	107
30. Working with NoSQL technologies	108
30.1. Redis	108
Connecting to Redis	108
30.2. MongoDB	108
Connecting to a MongoDB database	109
MongoTemplate	110
Spring Data MongoDB repositories	110
Embedded Mongo	110
30.3. Neo4j	111
Connecting to a Neo4j database	111
Using the embedded mode	111
Neo4jSession	112
Spring Data Neo4j repositories	112
Repository example	112
30.4. Gemfire	112
30.5. Solr	113

Connecting to Solr	113
Spring Data Solr repositories	113
30.6. Elasticsearch	113
Connecting to Elasticsearch using Jest	113
Connecting to Elasticsearch using Spring Data	114
Spring Data Elasticsearch repositories	114
30.7. Cassandra	114
Connecting to Cassandra	115
Spring Data Cassandra repositories	115
30.8. Couchbase	115
Connecting to Couchbase	115
Spring Data Couchbase repositories	116
30.9. LDAP	117
Connecting to an LDAP server	117
Spring Data LDAP repositories	117
Embedded in-memory LDAP server	117
30.10. InfluxDB	118
Connecting to InfluxDB	118
31. Caching	119
31.1. Supported cache providers	120
Generic	121
JCache (JSR-107)	121
EhCache 2.x	122
Hazelcast	122
Infinispan	122
Couchbase	122
Redis	123
Caffeine	123
Simple	123
None	124
32. Messaging	125
32.1. JMS	125
ActiveMQ support	125
Artemis support	126
Using a JNDI ConnectionFactory	126
Sending a message	126
Receiving a message	127
32.2. AMQP	128
RabbitMQ support	128
Sending a message	128
Receiving a message	129
32.3. Apache Kafka Support	130
Sending a Message	131
Receiving a Message	131
Additional Kafka Properties	131
33. Calling REST services with 'RestTemplate'	133
33.1. RestTemplate customization	133
34. Calling REST services with 'WebClient'	135
34.1. WebClient customization	135
35. Validation	136

36. Sending email	137
37. Distributed Transactions with JTA	138
37.1. Using an Atomikos transaction manager	138
37.2. Using a Bitronix transaction manager	138
37.3. Using a Narayana transaction manager	139
37.4. Using a Java EE managed transaction manager	139
37.5. Mixing XA and non-XA JMS connections	139
37.6. Supporting an alternative embedded transaction manager	140
38. Hazelcast	141
39. Quartz Scheduler	142
40. Spring Integration	143
41. Spring Session	144
42. Monitoring and management over JMX	145
43. Testing	146
43.1. Test scope dependencies	146
43.2. Testing Spring applications	146
43.3. Testing Spring Boot applications	146
Detecting test configuration	147
Excluding test configuration	148
Working with random ports	148
Mocking and spying beans	149
Auto-configured tests	150
Auto-configured JSON tests	150
Auto-configured Spring MVC tests	151
Auto-configured Spring WebFlux tests	153
Auto-configured Data JPA tests	154
Auto-configured JDBC tests	155
Auto-configured jOOQ tests	155
Auto-configured Data MongoDB tests	156
Auto-configured Data Neo4j tests	156
Auto-configured Data Redis tests	157
Auto-configured Data LDAP tests	157
Auto-configured REST clients	158
Auto-configured Spring REST Docs tests	158
Using Spock to test Spring Boot applications	159
43.4. Test utilities	160
ConfigFileApplicationContextInitializer	160
EnvironmentTestUtils	160
OutputCapture	160
TestRestTemplate	160
44. WebSockets	162
45. Web Services	163
46. Creating your own auto-configuration	164
46.1. Understanding auto-configured beans	164
46.2. Locating auto-configuration candidates	164
46.3. Condition annotations	164
Class conditions	165
Bean conditions	165
Property conditions	166
Resource conditions	166

Web application conditions	166
SpEL expression conditions	166
46.4. Creating your own starter	166
Naming	166
Autoconfigure module	167
Starter module	167
47. What to read next	168
V. Spring Boot Actuator: Production-ready features	169
48. Enabling production-ready features	170
49. Endpoints	171
49.1. Customizing endpoints	172
49.2. Hypermedia for actuator MVC endpoints	173
49.3. CORS support	173
49.4. Adding custom endpoints	173
49.5. Health information	174
49.6. Security with HealthIndicators	174
Auto-configured HealthIndicators	174
Writing custom HealthIndicators	175
49.7. Application information	176
Auto-configured InfoContributors	176
Custom application info information	176
Git commit information	177
Build information	177
Writing custom InfoContributors	177
50. Monitoring and management over HTTP	179
50.1. Accessing sensitive endpoints	179
50.2. Customizing the management endpoint paths	180
50.3. Customizing the management server port	180
50.4. Configuring management-specific SSL	181
50.5. Customizing the management server address	181
50.6. Disabling HTTP endpoints	181
50.7. HTTP health endpoint format and access restrictions	181
51. Monitoring and management over JMX	183
51.1. Customizing MBean names	183
51.2. Disabling JMX endpoints	183
51.3. Using Jolokia for JMX over HTTP	183
Customizing Jolokia	183
Disabling Jolokia	183
52. Loggers	185
52.1. Configure a Logger	185
53. Metrics	186
53.1. System metrics	186
53.2. DataSource metrics	187
53.3. Cache metrics	187
53.4. Tomcat session metrics	188
53.5. Recording your own metrics	188
53.6. Adding your own public metrics	188
53.7. Metric writers, exporters and aggregation	189
Example: Export to Redis	189
Example: Export to Open TSDB	190

Example: Export to Statsd	191
Example: Export to JMX	191
53.8. Aggregating metrics from multiple sources	191
53.9. Dropwizard Metrics	192
53.10. Message channel integration	192
54. Auditing	193
55. Tracing	194
55.1. Custom tracing	194
56. Process monitoring	195
56.1. Extend configuration	195
56.2. Programmatically	195
57. Cloud Foundry support	196
57.1. Disabling extended Cloud Foundry actuator support	196
57.2. Cloud Foundry self signed certificates	196
57.3. Custom security configuration	196
58. What to read next	197
VI. Deploying Spring Boot applications	198
59. Deploying to the cloud	199
59.1. Cloud Foundry	199
Binding to services	200
59.2. Heroku	201
59.3. OpenShift	202
59.4. Amazon Web Services (AWS)	202
AWS Elastic Beanstalk	203
Using the Tomcat platform	203
Using the Java SE platform	203
Best practices	203
Summary	203
59.5. Boxfuse and Amazon Web Services	203
59.6. Google Cloud	204
60. Installing Spring Boot applications	206
60.1. Supported operating systems	206
60.2. Unix/Linux services	206
Installation as an init.d service (System V)	206
Securing an init.d service	207
Installation as a systemd service	208
Customizing the startup script	208
Customizing script when it's written	209
Customizing script when it runs	210
60.3. Microsoft Windows services	211
61. What to read next	212
VII. Spring Boot CLI	213
62. Installing the CLI	214
63. Using the CLI	215
63.1. Running applications using the CLI	215
Deduced "grab" dependencies	216
Deduced "grab" coordinates	217
Default import statements	217
Automatic main method	217
Custom dependency management	217

63.2. Applications with multiple source files	218
63.3. Packaging your application	218
63.4. Initialize a new project	218
63.5. Using the embedded shell	219
63.6. Adding extensions to the CLI	219
64. Developing application with the Groovy beans DSL	220
65. Configuring the CLI with settings.xml	221
66. What to read next	222
VIII. Build tool plugins	223
67. Spring Boot Maven plugin	224
67.1. Including the plugin	224
67.2. Packaging executable jar and war files	225
68. Spring Boot Gradle plugin	226
69. Spring Boot AntLib module	227
69.1. Spring Boot Ant tasks	227
spring-boot:exejar	227
Examples	227
69.2. spring-boot:findmainclass	228
Examples	228
70. Supporting other build systems	229
70.1. Repackaging archives	229
70.2. Nested libraries	229
70.3. Finding a main class	229
70.4. Example repackage implementation	229
71. What to read next	230
IX. 'How-to' guides	231
72. Spring Boot application	232
72.1. Create your own FailureAnalyzer	232
72.2. Troubleshoot auto-configuration	232
72.3. Customize the Environment or ApplicationContext before it starts	233
72.4. Build an ApplicationContext hierarchy (adding a parent or root context)	234
72.5. Create a non-web application	234
73. Properties & configuration	235
73.1. Automatically expand properties at build time	235
Automatic property expansion using Maven	235
Automatic property expansion using Gradle	235
73.2. Externalize the configuration of SpringApplication	236
73.3. Change the location of external properties of an application	237
73.4. Use 'short' command line arguments	237
73.5. Use YAML for external properties	238
73.6. Set the active Spring profiles	238
73.7. Change configuration depending on the environment	238
73.8. Discover built-in options for external properties	239
74. Embedded Web servers	240
74.1. Use another Web server	240
74.2. Configure Jetty	241
74.3. Add a Servlet, Filter or Listener to an application	241
Add a Servlet, Filter or Listener using a Spring bean	241
Disable registration of a Servlet or Filter	241
Add Servlets, Filters, and Listeners using classpath scanning	242

74.4. Change the HTTP port	242
74.5. Use a random unassigned HTTP port	242
74.6. Discover the HTTP port at runtime	242
74.7. Configure SSL	243
74.8. Configure Access Logging	243
74.9. Use behind a front-end proxy server	243
Customize Tomcat's proxy configuration	244
74.10. Configure Tomcat	244
74.11. Enable Multiple Connectors with Tomcat	245
74.12. Use Tomcat's LegacyCookieProcessor	245
74.13. Configure Undertow	245
74.14. Enable Multiple Listeners with Undertow	246
74.15. Create WebSocket endpoints using @ServerEndpoint	246
74.16. Enable HTTP response compression	246
75. Spring MVC	247
75.1. Write a JSON REST service	247
75.2. Write an XML REST service	247
75.3. Customize the Jackson ObjectMapper	247
75.4. Customize the @ResponseBody rendering	249
75.5. Handling Multipart File Uploads	249
75.6. Switch off the Spring MVC DispatcherServlet	249
75.7. Switch off the Default MVC configuration	250
75.8. Customize ViewResolvers	250
76. HTTP clients	252
76.1. Configure RestTemplate to use a proxy	252
77. Logging	253
77.1. Configure Logback for logging	253
Configure logback for file only output	254
77.2. Configure Log4j for logging	254
Use YAML or JSON to configure Log4j 2	255
78. Data Access	256
78.1. Configure a custom DataSource	256
78.2. Configure Two DataSources	258
78.3. Use Spring Data repositories	259
78.4. Separate @Entity definitions from Spring configuration	259
78.5. Configure JPA properties	259
78.6. Configure Hibernate Naming Strategy	260
78.7. Use a custom EntityManagerFactory	260
78.8. Use Two EntityManagers	260
78.9. Use a traditional persistence.xml	261
78.10. Use Spring Data JPA and Mongo repositories	261
78.11. Expose Spring Data repositories as REST endpoint	262
78.12. Configure a component that is used by JPA	262
79. Database initialization	263
79.1. Initialize a database using JPA	263
79.2. Initialize a database using Hibernate	263
79.3. Initialize a database using Spring JDBC	263
79.4. Initialize a Spring Batch database	264
79.5. Use a higher-level database migration tool	264
Execute Flyway database migrations on startup	264

Execute Liquibase database migrations on startup	265
80. Messaging	266
80.1. Disable transacted JMS session	266
81. Batch applications	267
81.1. Execute Spring Batch jobs on startup	267
82. Actuator	268
82.1. Change the HTTP port or address of the actuator endpoints	268
82.2. Customize the 'whitelabel' error page	268
82.3. Actuator and Jersey	268
83. Security	269
83.1. Switch off the Spring Boot security configuration	269
83.2. Change the AuthenticationManager and add user accounts	269
83.3. Enable HTTPS when running behind a proxy server	269
84. Hot swapping	271
84.1. Reload static content	271
84.2. Reload templates without restarting the container	271
Thymeleaf templates	271
FreeMarker templates	271
Groovy templates	271
84.3. Fast application restarts	271
84.4. Reload Java classes without restarting the container	271
85. Build	272
85.1. Generate build information	272
85.2. Generate git information	272
85.3. Customize dependency versions	273
85.4. Create an executable JAR with Maven	273
85.5. Use a Spring Boot application as a dependency	274
85.6. Extract specific libraries when an executable jar runs	274
85.7. Create a non-executable JAR with exclusions	275
85.8. Remote debug a Spring Boot application started with Maven	275
85.9. Build an executable archive from Ant without using spring-boot-antlib	276
86. Traditional deployment	277
86.1. Create a deployable war file	277
86.2. Create a deployable war file for older servlet containers	278
86.3. Convert an existing application to Spring Boot	278
86.4. Deploying a WAR to WebLogic	280
86.5. Deploying a WAR in an Old (Servlet 2.5) Container	280
86.6. Use Lettuce instead of Jedis	281
X. Appendices	283
A. Common application properties	284
B. Configuration meta-data	307
B.1. Meta-data format	307
Group Attributes	308
Property Attributes	309
Hint Attributes	311
Repeated meta-data items	311
B.2. Providing manual hints	312
Value hint	312
Value provider	312
Any	313

Class reference	314
Handle As	314
Logger name	315
Spring bean reference	316
Spring profile name	317
B.3. Generating your own meta-data using the annotation processor	317
Nested properties	318
Adding additional meta-data	318
C. Auto-configuration classes	320
C.1. From the “spring-boot-autoconfigure” module	320
C.2. From the “spring-boot-actuator” module	323
D. Test auto-configuration annotations	325
E. The executable jar format	327
E.1. Nested JARs	327
The executable jar file structure	327
The executable war file structure	327
E.2. Spring Boot’s “JarFile” class	328
Compatibility with the standard Java “JarFile”	328
E.3. Launching executable jars	328
Launcher manifest	329
Exploded archives	329
E.4. PropertiesLauncher Features	329
E.5. Executable jar restrictions	331
Zip entry compression	331
System ClassLoader	331
E.6. Alternative single jar solutions	331
F. Dependency versions	332

Part I. Spring Boot Documentation

This section provides a brief overview of Spring Boot reference documentation. Think of it as map for the rest of the document. You can read this reference guide in a linear fashion, or you can skip sections if something doesn't interest you.

1. About the documentation

The Spring Boot reference guide is available as [html](#), [pdf](#) and [epub](#) documents. The latest copy is available at docs.spring.io/spring-boot/docs/current/reference.

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this Copyright Notice, whether distributed in print or electronically.

2. Getting help

Having trouble with Spring Boot, We'd like to help!

- Try the [How-to's](#) — they provide solutions to the most common questions.
- Learn the Spring basics — Spring Boot builds on many other Spring projects, check the [spring.io](#) website for a wealth of reference documentation. If you are just starting out with Spring, try one of the [guides](#).
- Ask a question - we monitor [stackoverflow.com](#) for questions tagged with [spring-boot](#).
- Report bugs with Spring Boot at [github.com/spring-projects/spring-boot/issues](#).

Note

All of Spring Boot is open source, including the documentation! If you find problems with the docs; or if you just want to improve them, please [get involved](#).

3. First steps

If you're just getting started with Spring Boot, or 'Spring' in general, [this is the place to start!](#)

- **From scratch:** [Overview](#) | [Requirements](#) | [Installation](#)
- **Tutorial:** [Part 1](#) | [Part 2](#)
- **Running your example:** [Part 1](#) | [Part 2](#)

4. Working with Spring Boot

Ready to actually start using Spring Boot? [We've got you covered.](#)

- **Build systems:** [Maven](#) | [Gradle](#) | [Ant](#) | [Starters](#)
- **Best practices:** [Code Structure](#) | [@Configuration](#) | [@EnableAutoConfiguration](#) | [Beans and Dependency Injection](#)
- **Running your code** [IDE](#) | [Packaged](#) | [Maven](#) | [Gradle](#)
- **Packaging your app:** [Production jars](#)
- **Spring Boot CLI:** [Using the CLI](#)

5. Learning about Spring Boot features

Need more details about Spring Boot's core features? [This is for you!](#)

- **Core Features:** [SpringApplication](#) | [External Configuration](#) | [Profiles](#) | [Logging](#)
- **Web Applications:** [MVC](#) | [Embedded Containers](#)
- **Working with data:** [SQL](#) | [NO-SQL](#)
- **Messaging:** [Overview](#) | [JMS](#)
- **Testing:** [Overview](#) | [Boot Applications](#) | [Utils](#)
- **Extending:** [Auto-configuration](#) | [@Conditions](#)

6. Moving to production

When you're ready to push your Spring Boot application to production, we've got [some tricks that you might like!](#)

- **Management endpoints:** [Overview](#) | [Customization](#)
- **Connection options:** [HTTP](#) | [JMX](#) |
- **Monitoring:** [Metrics](#) | [Auditing](#) | [Tracing](#) | [Process](#)

7. Advanced topics

Lastly, we have a few topics for the more advanced user.

- **Deploy Spring Boot Applications:** [Cloud Deployment](#) | [OS Service](#)
- **Build tool plugins:** [Maven](#) | [Gradle](#)
- **Appendix:** [Application Properties](#) | [Auto-configuration classes](#) | [Executable Jars](#)

Part II. Getting started

If you're just getting started with Spring Boot, or 'Spring' in general, this is the section for you! Here we answer the basic "what?", "how?" and "why?" questions. You'll find a gentle introduction to Spring Boot along with installation instructions. We'll then build our first Spring Boot application, discussing some core principles as we go.

8. Introducing Spring Boot

Spring Boot makes it easy to create stand-alone, production-grade Spring based Applications that you can “just run”. We take an opinionated view of the Spring platform and third-party libraries so you can get started with minimum fuss. Most Spring Boot applications need very little Spring configuration.

You can use Spring Boot to create Java applications that can be started using `java -jar` or more traditional war deployments. We also provide a command line tool that runs “spring scripts”.

Our primary goals are:

- Provide a radically faster and widely accessible getting started experience for all Spring development.
- Be opinionated out of the box, but get out of the way quickly as requirements start to diverge from the defaults.
- Provide a range of non-functional features that are common to large classes of projects (e.g. embedded servers, security, metrics, health checks, externalized configuration).
- Absolutely no code generation and no requirement for XML configuration.

9. System Requirements

Spring Boot 2.0.0.M3 requires [Java 8](#) and Spring Framework 5.0.0.RC3 or above. Explicit build support is provided for Maven (3.2+), and Gradle 3 (3.4 or later).

9.1 Servlet containers

The following embedded servlet containers are supported out of the box:

Name	Servlet Version
Tomcat 8.5	3.1
Jetty 9.4	3.1
Undertow 1.3	3.1

You can also deploy Spring Boot applications to any Servlet 3.0+ compatible container.

10. Installing Spring Boot

Spring Boot can be used with “classic” Java development tools or installed as a command line tool. Regardless, you will need [Java SDK v1.8](#) or higher. You should check your current Java installation before you begin:

```
$ java -version
```

If you are new to Java development, or if you just want to experiment with Spring Boot you might want to try the [Spring Boot CLI](#) first, otherwise, read on for “classic” installation instructions.

10.1 Installation instructions for the Java developer

You can use Spring Boot in the same way as any standard Java library. Simply include the appropriate `spring-boot-*.jar` files on your classpath. Spring Boot does not require any special tools integration, so you can use any IDE or text editor; and there is nothing special about a Spring Boot application, so you can run and debug as you would any other Java program.

Although you *could* just copy Spring Boot jars, we generally recommend that you use a build tool that supports dependency management (such as Maven or Gradle).

Maven installation

Spring Boot is compatible with Apache Maven 3.2 or above. If you don't already have Maven installed you can follow the instructions at maven.apache.org.

Tip

On many operating systems Maven can be installed via a package manager. If you're an OSX Homebrew user try `brew install maven`. Ubuntu users can run `sudo apt-get install maven`. Windows users with Chocolatey can run `choco install maven` from an elevated prompt.

Spring Boot dependencies use the `org.springframework.boot` groupId. Typically your Maven POM file will inherit from the `spring-boot-starter-parent` project and declare dependencies to one or more “[Starters](#)”. Spring Boot also provides an optional [Maven plugin](#) to create executable jars.

Here is a typical `pom.xml` file:

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>

  <groupId>com.example</groupId>
  <artifactId>myproject</artifactId>
  <version>0.0.1-SNAPSHOT</version>

  <!-- Inherit defaults from Spring Boot -->
  <parent>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-parent</artifactId>
    <version>2.0.0.M3</version>
  </parent>

  <!-- Add typical dependencies for a web application -->
  <dependencies>
```

```

    <dependency>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-starter-web</artifactId>
    </dependency>
  </dependencies>

  <!-- Package as an executable jar -->
  <build>
    <plugins>
      <plugin>
        <groupId>org.springframework.boot</groupId>
        <artifactId>spring-boot-maven-plugin</artifactId>
      </plugin>
    </plugins>
  </build>

  <!-- Add Spring repositories -->
  <!-- (you don't need this if you are using a .RELEASE version) -->
  <repositories>
    <repository>
      <id>spring-snapshots</id>
      <url>http://repo.spring.io/snapshot</url>
      <snapshots><enabled>true</enabled></snapshots>
    </repository>
    <repository>
      <id>spring-milestones</id>
      <url>http://repo.spring.io/milestone</url>
    </repository>
  </repositories>
  <pluginRepositories>
    <pluginRepository>
      <id>spring-snapshots</id>
      <url>http://repo.spring.io/snapshot</url>
    </pluginRepository>
    <pluginRepository>
      <id>spring-milestones</id>
      <url>http://repo.spring.io/milestone</url>
    </pluginRepository>
  </pluginRepositories>
</project>

```

Tip

The `spring-boot-starter-parent` is a great way to use Spring Boot, but it might not be suitable all of the time. Sometimes you may need to inherit from a different parent POM, or you might just not like our default settings. See [the section called “Using Spring Boot without the parent POM”](#) for an alternative solution that uses an `import` scope.

Gradle installation

Spring Boot is compatible with Gradle 3 (3.4 or later). If you don't already have Gradle installed you can follow the instructions at www.gradle.org/.

Spring Boot dependencies can be declared using the `org.springframework.boot` group. Typically your project will declare dependencies to one or more “[Starters](#)”. Spring Boot provides a useful [Gradle plugin](#) that can be used to simplify dependency declarations and to create executable jars.

Gradle Wrapper

The Gradle Wrapper provides a nice way of “obtaining” Gradle when you need to build a project. It's a small script and library that you commit alongside your code to bootstrap the build process. See docs.gradle.org/3.4.1/userguide/gradle_wrapper.html for details.

Here is a typical `build.gradle` file:

```
buildscript {
    repositories {
        jcenter()
        maven { url 'http://repo.spring.io/snapshot' }
        maven { url 'http://repo.spring.io/milestone' }
    }
    dependencies {
        classpath 'org.springframework.boot:spring-boot-gradle-plugin:2.0.0.M3'
    }
}

apply plugin: 'java'
apply plugin: 'org.springframework.boot'
apply plugin: 'io.spring.dependency-management'

jar {
    baseName = 'myproject'
    version = '0.0.1-SNAPSHOT'
}

repositories {
    jcenter()
    maven { url "http://repo.spring.io/snapshot" }
    maven { url "http://repo.spring.io/milestone" }
}

dependencies {
    compile("org.springframework.boot:spring-boot-starter-web")
    testCompile("org.springframework.boot:spring-boot-starter-test")
}
```

10.2 Installing the Spring Boot CLI

The Spring Boot CLI is a command line tool that can be used if you want to quickly prototype with Spring. It allows you to run [Groovy](#) scripts, which means that you have a familiar Java-like syntax, without so much boilerplate code.

You don't need to use the CLI to work with Spring Boot but it's definitely the quickest way to get a Spring application off the ground.

Manual installation

You can download the Spring CLI distribution from the Spring software repository:

- [spring-boot-cli-2.0.0.M3-bin.zip](#)
- [spring-boot-cli-2.0.0.M3-bin.tar.gz](#)

Cutting edge [snapshot distributions](#) are also available.

Once downloaded, follow the [INSTALL.txt](#) instructions from the unpacked archive. In summary: there is a `spring` script (`spring.bat` for Windows) in a `bin/` directory in the `.zip` file, or alternatively you can use `java -jar` with the `.jar` file (the script helps you to be sure that the classpath is set correctly).

Installation with SDKMAN!

SDKMAN! (The Software Development Kit Manager) can be used for managing multiple versions of various binary SDKs, including Groovy and the Spring Boot CLI. Get SDKMAN! from [sdkman.io](#) and install Spring Boot with

```
$ sdk install springboot
$ spring --version
Spring Boot v2.0.0.M3
```

If you are developing features for the CLI and want easy access to the version you just built, follow these extra instructions.

```
$ sdk install springboot dev /path/to/spring-boot/spring-boot-cli/target/spring-boot-cli-2.0.0.M3-bin/
spring-2.0.0.M3/
$ sdk default springboot dev
$ spring --version
Spring CLI v2.0.0.M3
```

This will install a local instance of `spring` called the `dev` instance. It points at your target build location, so every time you rebuild Spring Boot, `spring` will be up-to-date.

You can see it by doing this:

```
$ sdk ls springboot

=====
Available Springboot Versions
=====
> + dev
* 2.0.0.M3

=====

+ - local version
* - installed
> - currently in use
=====
```

OSX Homebrew installation

If you are on a Mac and using [Homebrew](#), all you need to do to install the Spring Boot CLI is:

```
$ brew tap pivotal/tap
$ brew install springboot
```

Homebrew will install `spring` to `/usr/local/bin`.

Note

If you don't see the formula, your installation of brew might be out-of-date. Just execute `brew update` and try again.

MacPorts installation

If you are on a Mac and using [MacPorts](#), all you need to do to install the Spring Boot CLI is:

```
$ sudo port install spring-boot-cli
```

Command-line completion

Spring Boot CLI ships with scripts that provide command completion for [BASH](#) and [zsh](#) shells. You can source the script (also named `spring`) in any shell, or put it in your personal or system-wide bash completion initialization. On a Debian system the system-wide scripts are in `/shell-completion/bash` and all scripts in that directory are executed when a new shell starts. To run the script manually, e.g. if you have installed using SDKMAN!

```
$ . ~/.sdkman/candidates/springboot/current/shell-completion/bash/spring
$ spring <HIT TAB HERE>
grab help jar run test version
```

Note

If you install Spring Boot CLI using Homebrew or MacPorts, the command-line completion scripts are automatically registered with your shell.

Quick start Spring CLI example

Here's a really simple web application that you can use to test your installation. Create a file called `app.groovy`:

```
@RestController
class ThisWillActuallyRun {

    @RequestMapping("/")
    String home() {
        "Hello World!"
    }
}
```

Then simply run it from a shell:

```
$ spring run app.groovy
```

Note

It will take some time when you first run the application as dependencies are downloaded. Subsequent runs will be much quicker.

Open localhost:8080 in your favorite web browser and you should see the following output:

```
Hello World!
```

10.3 Upgrading from an earlier version of Spring Boot

If you are upgrading from an earlier release of Spring Boot check the “release notes” hosted on the [project wiki](#). You'll find upgrade instructions along with a list of “new and noteworthy” features for each release.

To upgrade an existing CLI installation use the appropriate package manager command (for example `brew upgrade`) or, if you manually installed the CLI, follow the [standard instructions](#) remembering to update your `PATH` environment variable to remove any older references.

11. Developing your first Spring Boot application

Let's develop a simple "Hello World!" web application in Java that highlights some of Spring Boot's key features. We'll use Maven to build this project since most IDEs support it.

Tip

The spring.io web site contains many "Getting Started" guides that use Spring Boot. If you're looking to solve a specific problem; check there first.

You can shortcut the steps below by going to start.spring.io and choosing the web starter from the dependencies searcher. This will automatically generate a new project structure so that you can [start coding right away](#). Check the [documentation for more details](#).

Before we begin, open a terminal to check that you have valid versions of Java and Maven installed.

```
$ java -version
java version "1.8.0_102"
Java(TM) SE Runtime Environment (build 1.8.0_102-b14)
Java HotSpot(TM) 64-Bit Server VM (build 25.102-b14, mixed mode)
```

```
$ mvn -v
Apache Maven 3.3.9 (bb52d8502b132ec0a5a3f4c09453c07478323dc5; 2015-11-10T16:41:47+00:00)
Maven home: /usr/local/Cellar/maven/3.3.9/libexec
Java version: 1.8.0_102, vendor: Oracle Corporation
```

Note

This sample needs to be created in its own folder. Subsequent instructions assume that you have created a suitable folder and that it is your "current directory".

11.1 Creating the POM

We need to start by creating a Maven `pom.xml` file. The `pom.xml` is the recipe that will be used to build your project. Open your favorite text editor and add the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>

  <groupId>com.example</groupId>
  <artifactId>myproject</artifactId>
  <version>0.0.1-SNAPSHOT</version>

  <parent>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-parent</artifactId>
    <version>2.0.0.M3</version>
  </parent>

  <!-- Additional lines to be added here... -->

  <!-- (you don't need this if you are using a .RELEASE version) -->
  <repositories>
    <repository>
      <id>spring-snapshots</id>
      <url>http://repo.spring.io/snapshot</url>
      <snapshots><enabled>true</enabled></snapshots>
```

```

</repository>
<repository>
  <id>spring-milestones</id>
  <url>http://repo.spring.io/milestone</url>
</repository>
</repositories>
<pluginRepositories>
  <pluginRepository>
    <id>spring-snapshots</id>
    <url>http://repo.spring.io/snapshot</url>
  </pluginRepository>
  <pluginRepository>
    <id>spring-milestones</id>
    <url>http://repo.spring.io/milestone</url>
  </pluginRepository>
</pluginRepositories>
</project>

```

This should give you a working build, you can test it out by running `mvn package` (you can ignore the “jar will be empty - no content was marked for inclusion!” warning for now).

Note

At this point you could import the project into an IDE (most modern Java IDE’s include built-in support for Maven). For simplicity, we will continue to use a plain text editor for this example.

11.2 Adding classpath dependencies

Spring Boot provides a number of “Starters” that make easy to add jars to your classpath. Our sample application has already used `spring-boot-starter-parent` in the `parent` section of the POM. The `spring-boot-starter-parent` is a special starter that provides useful Maven defaults. It also provides a [dependency-management](#) section so that you can omit `version` tags for “blessed” dependencies.

Other “Starters” simply provide dependencies that you are likely to need when developing a specific type of application. Since we are developing a web application, we will add a `spring-boot-starter-web` dependency — but before that, let’s look at what we currently have.

```

$ mvn dependency:tree

[INFO] com.example:myproject:jar:0.0.1-SNAPSHOT

```

The `mvn dependency:tree` command prints a tree representation of your project dependencies. You can see that `spring-boot-starter-parent` provides no dependencies by itself. Let’s edit our `pom.xml` and add the `spring-boot-starter-web` dependency just below the `parent` section:

```

<dependencies>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-web</artifactId>
  </dependency>
</dependencies>

```

If you run `mvn dependency:tree` again, you will see that there are now a number of additional dependencies, including the Tomcat web server and Spring Boot itself.

11.3 Writing the code

To finish our application we need to create a single Java file. Maven will compile sources from `src/main/java` by default so you need to create that folder structure, then add a file named `src/main/java/Example.java`:

```
import org.springframework.boot.*;
import org.springframework.boot.autoconfigure.*;
import org.springframework.stereotype.*;
import org.springframework.web.bind.annotation.*;

@RestController
@EnableAutoConfiguration
public class Example {

    @RequestMapping("/")
    String home() {
        return "Hello World!";
    }

    public static void main(String[] args) throws Exception {
        SpringApplication.run(Example.class, args);
    }
}
```

Although there isn't much code here, quite a lot is going on. Let's step through the important parts.

The `@RestController` and `@RequestMapping` annotations

The first annotation on our `Example` class is `@RestController`. This is known as a *stereotype* annotation. It provides hints for people reading the code, and for Spring, that the class plays a specific role. In this case, our class is a web `@Controller` so Spring will consider it when handling incoming web requests.

The `@RequestMapping` annotation provides "routing" information. It is telling Spring that any HTTP request with the path "/" should be mapped to the `home` method. The `@RestController` annotation tells Spring to render the resulting string directly back to the caller.

Tip

The `@RestController` and `@RequestMapping` annotations are Spring MVC annotations (they are not specific to Spring Boot). See the [MVC section](#) in the Spring Reference Documentation for more details.

The `@EnableAutoConfiguration` annotation

The second class-level annotation is `@EnableAutoConfiguration`. This annotation tells Spring Boot to "guess" how you will want to configure Spring, based on the jar dependencies that you have added. Since `spring-boot-starter-web` added Tomcat and Spring MVC, the auto-configuration will assume that you are developing a web application and setup Spring accordingly.

Starters and Auto-Configuration

Auto-configuration is designed to work well with "Starters", but the two concepts are not directly tied. You are free to pick-and-choose jar dependencies outside of the starters and Spring Boot will still do its best to auto-configure your application.


```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
    </plugin>
  </plugins>
</build>

```

Note

The `spring-boot-starter-parent` POM includes `<executions>` configuration to bind the `repackage` goal. If you are not using the parent POM you will need to declare this configuration yourself. See the [plugin documentation](#) for details.

Save your `pom.xml` and run `mvn package` from the command line:

```

$ mvn package

[INFO] Scanning for projects...
[INFO]
[INFO] -----
[INFO] Building myproject 0.0.1-SNAPSHOT
[INFO] -----
[INFO] ....
[INFO] --- maven-jar-plugin:2.4:jar (default-jar) @ myproject ---
[INFO] Building jar: /Users/developer/example/spring-boot-example/target/myproject-0.0.1-SNAPSHOT.jar
[INFO]
[INFO] --- spring-boot-maven-plugin:2.0.0.M3:repackage (default) @ myproject ---
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----

```

If you look in the `target` directory you should see `myproject-0.0.1-SNAPSHOT.jar`. The file should be around 10 MB in size. If you want to peek inside, you can use `jar tvf`:

```
$ jar tvf target/myproject-0.0.1-SNAPSHOT.jar
```

You should also see a much smaller file named `myproject-0.0.1-SNAPSHOT.jar.original` in the `target` directory. This is the original jar file that Maven created before it was repackaged by Spring Boot.

To run that application, use the `java -jar` command:

```

$ java -jar target/myproject-0.0.1-SNAPSHOT.jar

 . ____ _ _   _
 /\\ / ___'_ __ _ ( '_ ) ___ __ / \\  \\
( ( )\___| ' | ' _ \| | | |__/_ \| / ___ \|
 \\/ ___)| |_)| | | | |___/_| |_/ ___/| |_) |
  ' |___| |_)| | | | |___/_| |_/ ___/| |_) |
 =====|_|=====|__/=/_/___/
 :: Spring Boot :: (v2.0.0.M3)

.....
..... (log output here)
.....
..... Started Example in 2.536 seconds (JVM running for 2.864)

```

As before, to gracefully exit the application hit `ctrl-c`.

12. What to read next

Hopefully this section has provided you with some of the Spring Boot basics, and got you on your way to writing your own applications. If you're a task-oriented type of developer you might want to jump over to spring.io and check out some of the [getting started](#) guides that solve specific "How do I do that with Spring" problems; we also have Spring Boot-specific [How-to](#) reference documentation.

The [Spring Boot repository](#) has also a [bunch of samples](#) you can run. The samples are independent of the rest of the code (that is you don't need to build the rest to run or use the samples).

Otherwise, the next logical step is to read [Part III, "Using Spring Boot"](#). If you're really impatient, you could also jump ahead and read about [Spring Boot features](#).

Part III. Using Spring Boot

This section goes into more detail about how you should use Spring Boot. It covers topics such as build systems, auto-configuration and how to run your applications. We also cover some Spring Boot best practices. Although there is nothing particularly special about Spring Boot (it is just another library that you can consume), there are a few recommendations that, when followed, will make your development process just a little easier.

If you're just starting out with Spring Boot, you should probably read the [Getting Started](#) guide before diving into this section.

13. Build systems

It is strongly recommended that you choose a build system that supports [dependency management](#), and one that can consume artifacts published to the “Maven Central” repository. We would recommend that you choose Maven or Gradle. It is possible to get Spring Boot to work with other build systems (Ant for example), but they will not be particularly well supported.

13.1 Dependency management

Each release of Spring Boot provides a curated list of dependencies it supports. In practice, you do not need to provide a version for any of these dependencies in your build configuration as Spring Boot is managing that for you. When you upgrade Spring Boot itself, these dependencies will be upgraded as well in a consistent way.

Note

You can still specify a version and override Spring Boot’s recommendations if you feel that’s necessary.

The curated list contains all the spring modules that you can use with Spring Boot as well as a refined list of third party libraries. The list is available as a standard [Bills of Materials \(spring-boot-dependencies\)](#) that can be used with both [Maven](#) and [Gradle](#).

Warning

Each release of Spring Boot is associated with a base version of the Spring Framework so we **highly** recommend you to not specify its version on your own.

13.2 Maven

Maven users can inherit from the `spring-boot-starter-parent` project to obtain sensible defaults. The parent project provides the following features:

- Java 1.8 as the default compiler level.
- UTF-8 source encoding.
- A [Dependency Management section](#), allowing you to omit `<version>` tags for common dependencies, inherited from the `spring-boot-dependencies` POM.
- Sensible [resource filtering](#).
- Sensible plugin configuration ([exec plugin](#), [surefire](#), [Git commit ID](#), [shade](#)).
- Sensible resource filtering for `application.properties` and `application.yml` including profile-specific files (e.g. `application-foo.properties` and `application-foo.yml`)

On the last point: since the default config files accept Spring style placeholders (`${...}`) the Maven filtering is changed to use `@..@` placeholders (you can override that with a Maven property `resource.delimiter`).

Inheriting the starter parent

To configure your project to inherit from the `spring-boot-starter-parent` simply set the parent:

```

<!-- Inherit defaults from Spring Boot -->
<parent>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-parent</artifactId>
  <version>2.0.0.M3</version>
</parent>

```

Note

You should only need to specify the Spring Boot version number on this dependency. If you import additional starters, you can safely omit the version number.

With that setup, you can also override individual dependencies by overriding a property in your own project. For instance, to upgrade to another Spring Data release train you'd add the following to your `pom.xml`.

```

<properties>
  <spring-data-releasetrain.version>Fowler-SR2</spring-data-releasetrain.version>
</properties>

```

Tip

Check the [spring-boot-dependencies pom](#) for a list of supported properties.

Using Spring Boot without the parent POM

Not everyone likes inheriting from the `spring-boot-starter-parent` POM. You may have your own corporate standard parent that you need to use, or you may just prefer to explicitly declare all your Maven configuration.

If you don't want to use the `spring-boot-starter-parent`, you can still keep the benefit of the dependency management (but not the plugin management) by using a `scope=import` dependency:

```

<dependencyManagement>
  <dependencies>
    <dependency>
      <!-- Import dependency management from Spring Boot -->
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-dependencies</artifactId>
      <version>2.0.0.M3</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>

```

That setup does not allow you to override individual dependencies using a property as explained above. To achieve the same result, you'd need to add an entry in the `dependencyManagement` of your project **before** the `spring-boot-dependencies` entry. For instance, to upgrade to another Spring Data release train you'd add the following to your `pom.xml`.

```

<dependencyManagement>
  <dependencies>
    <!-- Override Spring Data release train provided by Spring Boot -->
    <dependency>
      <groupId>org.springframework.data</groupId>
      <artifactId>spring-data-releasetrain</artifactId>
      <version>Fowler-SR2</version>
      <scope>import</scope>
      <type>pom</type>

```

```

</dependency>
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-dependencies</artifactId>
  <version>2.0.0.M3</version>
  <type>pom</type>
  <scope>import</scope>
</dependency>
</dependencies>
</dependencyManagement>

```

Note

In the example above, we specify a *BOM* but any dependency type can be overridden that way.

Using the Spring Boot Maven plugin

Spring Boot includes a [Maven plugin](#) that can package the project as an executable jar. Add the plugin to your `<plugins>` section if you want to use it:

```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
    </plugin>
  </plugins>
</build>

```

Note

If you use the Spring Boot starter parent pom, you only need to add the plugin, there is no need for to configure it unless you want to change the settings defined in the parent.

13.3 Gradle

To learn about using Spring Boot with Gradle, please refer to the documentation for Spring Boot's Gradle plugin:

- Reference ([HTML](#) and [PDF](#))
- [API](#)

13.4 Ant

It is possible to build a Spring Boot project using Apache Ant+Ivy. The `spring-boot-antlib` "AntLib" module is also available to help Ant create executable jars.

To declare dependencies a typical `ivy.xml` file will look something like this:

```

<ivy-module version="2.0">
  <info organisation="org.springframework.boot" module="spring-boot-sample-ant" />
  <configurations>
    <conf name="compile" description="everything needed to compile this module" />
    <conf name="runtime" extends="compile" description="everything needed to run this module" />
  </configurations>
  <dependencies>
    <dependency org="org.springframework.boot" name="spring-boot-starter"
      rev="{spring-boot.version}" conf="compile" />
  </dependencies>

```

```
</ivy-module>
```

A typical `build.xml` will look like this:

```
<project
  xmlns:ivy="antlib:org.apache.ivy.ant"
  xmlns:spring-boot="antlib:org.springframework.boot.ant"
  name="myapp" default="build">

  <property name="spring-boot.version" value="1.3.0.BUILD-SNAPSHOT" />

  <target name="resolve" description="--> retrieve dependencies with ivy">
    <ivy:retrieve pattern="lib/[conf]/[artifact]-[type]-[revision].[ext]" />
  </target>

  <target name="classpaths" depends="resolve">
    <path id="compile.classpath">
      <fileset dir="lib/compile" includes="*.jar" />
    </path>
  </target>

  <target name="init" depends="classpaths">
    <mkdir dir="build/classes" />
  </target>

  <target name="compile" depends="init" description="compile">
    <javac srcdir="src/main/java" destdir="build/classes" classpathref="compile.classpath" />
  </target>

  <target name="build" depends="compile">
    <spring-boot:exejar destfile="build/myapp.jar" classes="build/classes">
      <spring-boot:lib>
        <fileset dir="lib/runtime" />
      </spring-boot:lib>
    </spring-boot:exejar>
  </target>
</project>
```

Tip

See the [Section 85.9, “Build an executable archive from Ant without using `spring-boot-antlib`”](#) “How-to” if you don’t want to use the `spring-boot-antlib` module.

13.5 Starters

Starters are a set of convenient dependency descriptors that you can include in your application. You get a one-stop-shop for all the Spring and related technology that you need, without having to hunt through sample code and copy paste loads of dependency descriptors. For example, if you want to get started using Spring and JPA for database access, just include the `spring-boot-starter-data-jpa` dependency in your project, and you are good to go.

The starters contain a lot of the dependencies that you need to get a project up and running quickly and with a consistent, supported set of managed transitive dependencies.

What’s in a name

All **official** starters follow a similar naming pattern; `spring-boot-starter-*`, where `*` is a particular type of application. This naming structure is intended to help when you need to find a starter. The Maven integration in many IDEs allow you to search dependencies by name. For example, with the appropriate Eclipse or STS plugin installed, you can simply hit `ctrl-space` in the POM editor and type “spring-boot-starter” for a complete list.

As explained in the [Creating your own starter](#) section, third party starters should not start with `spring-boot` as it is reserved for official Spring Boot artifacts. A third-party starter for `acme` will be typically named `acme-spring-boot-starter`.

The following application starters are provided by Spring Boot under the `org.springframework.boot` group:

Table 13.1. Spring Boot application starters

Name	Description	Pom
<code>spring-boot-starter</code>	Core starter, including auto-configuration support, logging and YAML	Pom
<code>spring-boot-starter-activemq</code>	Starter for JMS messaging using Apache ActiveMQ	Pom
<code>spring-boot-starter-amqp</code>	Starter for using Spring AMQP and Rabbit MQ	Pom
<code>spring-boot-starter-aop</code>	Starter for aspect-oriented programming with Spring AOP and AspectJ	Pom
<code>spring-boot-starter-artemis</code>	Starter for JMS messaging using Apache Artemis	Pom
<code>spring-boot-starter-batch</code>	Starter for using Spring Batch	Pom
<code>spring-boot-starter-cache</code>	Starter for using Spring Framework's caching support	Pom
<code>spring-boot-starter-cloud-connectors</code>	Starter for using Spring Cloud Connectors which simplifies connecting to services in cloud platforms like Cloud Foundry and Heroku	Pom
<code>spring-boot-starter-data-cassandra</code>	Starter for using Cassandra distributed database and Spring Data Cassandra	Pom
<code>spring-boot-starter-data-cassandra-reactive</code>	Starter for using Cassandra distributed database and Spring Data Cassandra Reactive	Pom
<code>spring-boot-starter-data-couchbase</code>	Starter for using Couchbase document-oriented database and Spring Data Couchbase	Pom
<code>spring-boot-starter-data-elasticsearch</code>	Starter for using Elasticsearch search and analytics engine and Spring Data Elasticsearch	Pom

Name	Description	Pom
spring-boot-starter-data-jpa	Starter for using Spring Data JPA with Hibernate	Pom
spring-boot-starter-data-ldap	Starter for using Spring Data LDAP	Pom
spring-boot-starter-data-mongodb	Starter for using MongoDB document-oriented database and Spring Data MongoDB	Pom
spring-boot-starter-data-mongodb-reactive	Starter for using MongoDB document-oriented database and Spring Data MongoDB Reactive	Pom
spring-boot-starter-data-neo4j	Starter for using Neo4j graph database and Spring Data Neo4j	Pom
spring-boot-starter-data-redis	Starter for using Redis key-value data store with Spring Data Redis and the Jedis client	Pom
spring-boot-starter-data-redis-reactive	Starter for using Redis key-value data store with Spring Data Redis reactive and the Lettuce client	Pom
spring-boot-starter-data-rest	Starter for exposing Spring Data repositories over REST using Spring Data REST	Pom
spring-boot-starter-data-solr	Starter for using the Apache Solr search platform with Spring Data Solr	Pom
spring-boot-starter-freemarker	Starter for building MVC web applications using FreeMarker views	Pom
spring-boot-starter-groovy-templates	Starter for building MVC web applications using Groovy Templates views	Pom
spring-boot-starter-hateoas	Starter for building hypermedia-based RESTful web application with Spring MVC and Spring HATEOAS	Pom
spring-boot-starter-integration	Starter for using Spring Integration	Pom

Name	Description	Pom
spring-boot-starter-jdbc	Starter for using JDBC with the Tomcat JDBC connection pool	Pom
spring-boot-starter-jersey	Starter for building RESTful web applications using JAX-RS and Jersey. An alternative to spring-boot-starter-web	Pom
spring-boot-starter-jooq	Starter for using jOOQ to access SQL databases. An alternative to spring-boot-starter-data-jpa or spring-boot-starter-jdbc	Pom
spring-boot-starter-jta-atomikos	Starter for JTA transactions using Atomikos	Pom
spring-boot-starter-jta-bitronix	Starter for JTA transactions using Bitronix	Pom
spring-boot-starter-jta-narayana	Spring Boot Narayana JTA Starter	Pom
spring-boot-starter-mail	Starter for using Java Mail and Spring Framework's email sending support	Pom
spring-boot-starter-mobile	Starter for building web applications using Spring Mobile	Pom
spring-boot-starter-mustache	Starter for building web applications using Mustache views	Pom
spring-boot-starter-quartz	Spring Boot Quartz Starter	Pom
spring-boot-starter-security	Starter for using Spring Security	Pom
spring-boot-starter-social-facebook	Starter for using Spring Social Facebook	Pom
spring-boot-starter-social-linkedin	Stater for using Spring Social LinkedIn	Pom
spring-boot-starter-social-twitter	Starter for using Spring Social Twitter	Pom
spring-boot-starter-test	Starter for testing Spring Boot applications with libraries	Pom

Name	Description	Pom
	including JUnit, Hamcrest and Mockito	
spring-boot-starter-thymeleaf	Starter for building MVC web applications using Thymeleaf views	Pom
spring-boot-starter-validation	Starter for using Java Bean Validation with Hibernate Validator	Pom
spring-boot-starter-web	Starter for building web, including RESTful, applications using Spring MVC. Uses Tomcat as the default embedded container	Pom
spring-boot-starter-web-services	Starter for using Spring Web Services	Pom
spring-boot-starter-webflux	Starter for building WebFlux applications using Spring Framework's Reactive Web support	Pom
spring-boot-starter-websocket	Starter for building WebSocket applications using Spring Framework's WebSocket support	Pom

In addition to the application starters, the following starters can be used to add *production ready* features:

Table 13.2. Spring Boot production starters

Name	Description	Pom
spring-boot-starter-actuator	Starter for using Spring Boot's Actuator which provides production ready features to help you monitor and manage your application	Pom

Finally, Spring Boot also includes some starters that can be used if you want to exclude or swap specific technical facets:

Table 13.3. Spring Boot technical starters

Name	Description	Pom
spring-boot-starter-jetty	Starter for using Jetty as the embedded servlet container. An alternative to spring-boot-starter-tomcat	Pom

Name	Description	Pom
spring-boot-starter-json	Starter for reading and writing json	Pom
spring-boot-starter-log4j2	Starter for using Log4j2 for logging. An alternative to spring-boot-starter-logging	Pom
spring-boot-starter-logging	Starter for logging using Logback. Default logging starter	Pom
spring-boot-starter-reactor-netty	Starter for using Reactor Netty as the embedded reactive HTTP server.	Pom
spring-boot-starter-tomcat	Starter for using Tomcat as the embedded servlet container. Default servlet container starter used by spring-boot-starter-web	Pom
spring-boot-starter-undertow	Starter for using Undertow as the embedded servlet container. An alternative to spring-boot-starter-tomcat	Pom

Tip

For a list of additional community contributed starters, see the [README file](#) in the `spring-boot-starters` module on GitHub.

14. Structuring your code

Spring Boot does not require any specific code layout to work, however, there are some best practices that help.

14.1 Using the “default” package

When a class doesn't include a `package` declaration it is considered to be in the “default package”. The use of the “default package” is generally discouraged, and should be avoided. It can cause particular problems for Spring Boot applications that use `@ComponentScan`, `@EntityScan` or `@SpringBootApplication` annotations, since every class from every jar, will be read.

Tip

We recommend that you follow Java's recommended package naming conventions and use a reversed domain name (for example, `com.example.project`).

14.2 Locating the main application class

We generally recommend that you locate your main application class in a root package above other classes. The `@EnableAutoConfiguration` annotation is often placed on your main class, and it implicitly defines a base “search package” for certain items. For example, if you are writing a JPA application, the package of the `@EnableAutoConfiguration` annotated class will be used to search for `@Entity` items.

Using a root package also allows the `@ComponentScan` annotation to be used without needing to specify a `basePackage` attribute. You can also use the `@SpringBootApplication` annotation if your main class is in the root package.

Here is a typical layout:

```
com
+- example
  +- myproject
    +- Application.java
    |
    +- domain
    |   +- Customer.java
    |   +- CustomerRepository.java
    |
    +- service
    |   +- CustomerService.java
    |
    +- web
    |   +- CustomerController.java
```

The `Application.java` file would declare the main method, along with the basic `@Configuration`.

```
package com.example.myproject;

import org.springframework.boot.SpringApplication;
import org.springframework.boot.autoconfigure.EnableAutoConfiguration;
import org.springframework.context.annotation.ComponentScan;
import org.springframework.context.annotation.Configuration;

@Configuration
@EnableAutoConfiguration
@ComponentScan
```

```
public class Application {  
  
    public static void main(String[] args) {  
        SpringApplication.run(Application.class, args);  
    }  
  
}
```

15. Configuration classes

Spring Boot favors Java-based configuration. Although it is possible to use `SpringApplication` with an XML sources, we generally recommend that your primary source is a single `@Configuration` class. Usually the class that defines the `main` method is also a good candidate as the primary `@Configuration`.

Tip

Many Spring configuration examples have been published on the Internet that use XML configuration. Always try to use the equivalent Java-based configuration if possible. Searching for `Enable*` annotations can be a good starting point.

15.1 Importing additional configuration classes

You don't need to put all your `@Configuration` into a single class. The `@Import` annotation can be used to import additional configuration classes. Alternatively, you can use `@ComponentScan` to automatically pick up all Spring components, including `@Configuration` classes.

15.2 Importing XML configuration

If you absolutely must use XML based configuration, we recommend that you still start with a `@Configuration` class. You can then use an additional `@ImportResource` annotation to load XML configuration files.

16. Auto-configuration

Spring Boot auto-configuration attempts to automatically configure your Spring application based on the jar dependencies that you have added. For example, if `HSQLDB` is on your classpath, and you have not manually configured any database connection beans, then we will auto-configure an in-memory database.

You need to opt-in to auto-configuration by adding the `@EnableAutoConfiguration` or `@SpringBootApplication` annotations to one of your `@Configuration` classes.

Tip

You should only ever add one `@EnableAutoConfiguration` annotation. We generally recommend that you add it to your primary `@Configuration` class.

16.1 Gradually replacing auto-configuration

Auto-configuration is noninvasive, at any point you can start to define your own configuration to replace specific parts of the auto-configuration. For example, if you add your own `DataSource` bean, the default embedded database support will back away.

If you need to find out what auto-configuration is currently being applied, and why, start your application with the `--debug` switch. This will enable debug logs for a selection of core loggers and log an auto-configuration report to the console.

16.2 Disabling specific auto-configuration

If you find that specific auto-configure classes are being applied that you don't want, you can use the `exclude` attribute of `@EnableAutoConfiguration` to disable them.

```
import org.springframework.boot.autoconfigure.*;
import org.springframework.boot.autoconfigure.jdbc.*;
import org.springframework.context.annotation.*;

@Configuration
@EnableAutoConfiguration(exclude={DataSourceAutoConfiguration.class})
public class MyConfiguration {
}
```

If the class is not on the classpath, you can use the `excludeName` attribute of the annotation and specify the fully qualified name instead. Finally, you can also control the list of auto-configuration classes to exclude via the `spring.autoconfigure.exclude` property.

Tip

You can define exclusions both at the annotation level and using the property.

17. Spring Beans and dependency injection

You are free to use any of the standard Spring Framework techniques to define your beans and their injected dependencies. For simplicity, we often find that using `@ComponentScan` to find your beans, in combination with `@Autowired` constructor injection works well.

If you structure your code as suggested above (locating your application class in a root package), you can add `@ComponentScan` without any arguments. All of your application components (`@Component`, `@Service`, `@Repository`, `@Controller` etc.) will be automatically registered as Spring Beans.

Here is an example `@Service` Bean that uses constructor injection to obtain a required `RiskAssessor` bean.

```
package com.example.service;

import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.stereotype.Service;

@Service
public class DatabaseAccountService implements AccountService {

    private final RiskAssessor riskAssessor;

    @Autowired
    public DatabaseAccountService(RiskAssessor riskAssessor) {
        this.riskAssessor = riskAssessor;
    }

    // ...
}
```

And if a bean has one constructor, you can omit the `@Autowired`.

```
@Service
public class DatabaseAccountService implements AccountService {

    private final RiskAssessor riskAssessor;

    public DatabaseAccountService(RiskAssessor riskAssessor) {
        this.riskAssessor = riskAssessor;
    }

    // ...
}
```

Tip

Notice how using constructor injection allows the `riskAssessor` field to be marked as `final`, indicating that it cannot be subsequently changed.

18. Using the `@SpringBootApplication` annotation

Many Spring Boot developers always have their main class annotated with `@Configuration`, `@EnableAutoConfiguration` and `@ComponentScan`. Since these annotations are so frequently used together (especially if you follow the [best practices](#) above), Spring Boot provides a convenient `@SpringBootApplication` alternative.

The `@SpringBootApplication` annotation is equivalent to using `@Configuration`, `@EnableAutoConfiguration` and `@ComponentScan` with their default attributes:

```
package com.example.myproject;

import org.springframework.boot.SpringApplication;
import org.springframework.boot.autoconfigure.SpringBootApplication;

@SpringBootApplication // same as @Configuration @EnableAutoConfiguration @ComponentScan
public class Application {

    public static void main(String[] args) {
        SpringApplication.run(Application.class, args);
    }

}
```

Note

`@SpringBootApplication` also provides aliases to customize the attributes of `@EnableAutoConfiguration` and `@ComponentScan`.

19. Running your application

One of the biggest advantages of packaging your application as jar and using an embedded HTTP server is that you can run your application as you would any other. Debugging Spring Boot applications is also easy; you don't need any special IDE plugins or extensions.

Note

This section only covers jar based packaging, If you choose to package your application as a war file you should refer to your server and IDE documentation.

19.1 Running from an IDE

You can run a Spring Boot application from your IDE as a simple Java application, however, first you will need to import your project. Import steps will vary depending on your IDE and build system. Most IDEs can import Maven projects directly, for example Eclipse users can select `Import... → Existing Maven Projects` from the `File` menu.

If you can't directly import your project into your IDE, you may be able to generate IDE metadata using a build plugin. Maven includes plugins for [Eclipse](#) and [IDEA](#); Gradle offers plugins for [various IDEs](#).

Tip

If you accidentally run a web application twice you will see a "Port already in use" error. STS users can use the `Relaunch` button rather than `Run` to ensure that any existing instance is closed.

19.2 Running as a packaged application

If you use the Spring Boot Maven or Gradle plugins to create an executable jar you can run your application using `java -jar`. For example:

```
$ java -jar target/myproject-0.0.1-SNAPSHOT.jar
```

It is also possible to run a packaged application with remote debugging support enabled. This allows you to attach a debugger to your packaged application:

```
$ java -Xdebug -Xrunjdpw:server=y,transport=dt_socket,address=8000,suspend=n \  
-jar target/myproject-0.0.1-SNAPSHOT.jar
```

19.3 Using the Maven plugin

The Spring Boot Maven plugin includes a `run` goal which can be used to quickly compile and run your application. Applications run in an exploded form just like in your IDE.

```
$ mvn spring-boot:run
```

You might also want to use the useful operating system environment variable:

```
$ export MAVEN_OPTS=-Xmx1024m
```

19.4 Using the Gradle plugin

The Spring Boot Gradle plugin also includes a `bootRun` task which can be used to run your application in an exploded form. The `bootRun` task is added whenever you apply the `org.springframework.boot` and `java` plugins:

```
$ gradle bootRun
```

You might also want to use this useful operating system environment variable:

```
$ export JAVA_OPTS=-Xmx1024m
```

19.5 Hot swapping

Since Spring Boot applications are just plain Java applications, JVM hot-swapping should work out of the box. JVM hot swapping is somewhat limited with the bytecode that it can replace, for a more complete solution [JRebel](#) can be used. The `spring-boot-devtools` module also includes support for quick application restarts.

See the [Chapter 20, *Developer tools*](#) section below and the [Hot swapping “How-to”](#) for details.

20. Developer tools

Spring Boot includes an additional set of tools that can make the application development experience a little more pleasant. The `spring-boot-devtools` module can be included in any project to provide additional development-time features. To include devtools support, simply add the module dependency to your build:

Maven.

```
<dependencies>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-devtools</artifactId>
    <optional>true</optional>
  </dependency>
</dependencies>
```

Gradle.

```
dependencies {
    compile("org.springframework.boot:spring-boot-devtools")
}
```

Note

Developer tools are automatically disabled when running a fully packaged application. If your application is launched using `java -jar` or if it's started using a special classloader, then it is considered a "production application". Flagging the dependency as optional is a best practice that prevents devtools from being transitively applied to other modules using your project. Gradle does not support optional dependencies out-of-the-box so you may want to have a look to the [propdeps-plugin](#) in the meantime.

Tip

repackaged archives do not contain devtools by default. If you want to use [certain remote devtools feature](#), you'll need to disable the `excludeDevtools` build property to include it. The property is supported with both the Maven and Gradle plugins.

20.1 Property defaults

Several of the libraries supported by Spring Boot use caches to improve performance. For example, [template engines](#) will cache compiled templates to avoid repeatedly parsing template files. Also, Spring MVC can add HTTP caching headers to responses when serving static resources.

Whilst caching is very beneficial in production, it can be counter productive during development, preventing you from seeing the changes you just made in your application. For this reason, `spring-boot-devtools` will disable those caching options by default.

Cache options are usually configured by settings in your `application.properties` file. For example, Thymeleaf offers the `spring.thymeleaf.cache` property. Rather than needing to set these properties manually, the `spring-boot-devtools` module will automatically apply sensible development-time configuration.

Tip

For a complete list of the properties that are applied see [DevToolsPropertyDefaultsPostProcessor](#).

20.2 Automatic restart

Applications that use `spring-boot-devtools` will automatically restart whenever files on the classpath change. This can be a useful feature when working in an IDE as it gives a very fast feedback loop for code changes. By default, any entry on the classpath that points to a folder will be monitored for changes. Note that certain resources such as static assets and view templates [do not need to restart the application](#).

Triggering a restart

As DevTools monitors classpath resources, the only way to trigger a restart is to update the classpath. The way in which you cause the classpath to be updated depends on the IDE that you are using. In Eclipse, saving a modified file will cause the classpath to be updated and trigger a restart. In IntelliJ IDEA, building the project (`Build -> Make Project`) will have the same effect.

Note

You can also start your application via the supported build plugins (i.e. Maven and Gradle) as long as forking is enabled since DevTools need an isolated application classloader to operate properly. Gradle and Maven do that by default when they detect DevTools on the classpath.

Tip

Automatic restart works very well when used with LiveReload. [See below](#) for details. If you use JRebel automatic restarts will be disabled in favor of dynamic class reloading. Other devtools features (such as LiveReload and property overrides) can still be used.

Note

DevTools relies on the application context's shutdown hook to close it during a restart. It will not work correctly if you have disabled the shutdown hook (`SpringApplication.setRegisterShutdownHook(false)`).

Note

When deciding if an entry on the classpath should trigger a restart when it changes, DevTools automatically ignores projects named `spring-boot`, `spring-boot-devtools`, `spring-boot-autoconfigure`, `spring-boot-actuator`, and `spring-boot-starter`.

Note

DevTools needs to customize the `ResourceLoader` used by the `ApplicationContext`: if your application provides one already, it is going to be wrapped. Direct override of the `getResource` method on the `ApplicationContext` is not supported.

Restart vs Reload

The restart technology provided by Spring Boot works by using two classloaders. Classes that don't change (for example, those from third-party jars) are loaded into a *base* classloader. Classes that you're actively developing are loaded into a *restart* classloader. When the application is restarted, the *restart* classloader is thrown away and a new one is created. This approach means that application restarts are typically much faster than "cold starts" since the *base* classloader is already available and populated.

If you find that restarts aren't quick enough for your applications, or you encounter classloading issues, you could consider reloading technologies such as [JRebel](#) from ZeroTurnaround. These work by rewriting classes as they are loaded to make them more amenable to reloading.

Excluding resources

Certain resources don't necessarily need to trigger a restart when they are changed. For example, Thymeleaf templates can just be edited in-place. By default changing resources in `/META-INF/maven`, `/META-INF/resources`, `/resources`, `/static`, `/public` or `/templates` will not trigger a restart but will trigger a [live reload](#). If you want to customize these exclusions you can use the `spring.devtools.restart.exclude` property. For example, to exclude only `/static` and `/public` you would set the following:

```
spring.devtools.restart.exclude=static/**,public/**
```

Tip

if you want to keep those defaults and *add* additional exclusions, use the `spring.devtools.restart.additional-exclude` property instead.

Watching additional paths

You may want your application to be restarted or reloaded when you make changes to files that are not on the classpath. To do so, use the `spring.devtools.restart.additional-paths` property to configure additional paths to watch for changes. You can use the `spring.devtools.restart.exclude` property [described above](#) to control whether changes beneath the additional paths will trigger a full restart or just a [live reload](#).

Disabling restart

If you don't want to use the restart feature you can disable it using the `spring.devtools.restart.enabled` property. In most cases you can set this in your `application.properties` (this will still initialize the restart classloader but it won't watch for file changes).

If you need to *completely* disable restart support, for example, because it doesn't work with a specific library, you need to set a `System` property before calling `SpringApplication.run(...)`. For example:

```
public static void main(String[] args) {
    System.setProperty("spring.devtools.restart.enabled", "false");
    SpringApplication.run(MyApp.class, args);
}
```


Using a trigger file

If you work with an IDE that continuously compiles changed files, you might prefer to trigger restarts only at specific times. To do this you can use a “trigger file”, which is a special file that must be modified when you want to actually trigger a restart check. Changing the file only triggers the check and the restart will only occur if Devtools has detected it has to do something. The trigger file could be updated manually, or via an IDE plugin.

To use a trigger file use the `spring.devtools.restart.trigger-file` property.

Tip

You might want to set `spring.devtools.restart.trigger-file` as a [global setting](#) so that all your projects behave in the same way.

Customizing the restart classloader

As described in the [Restart vs Reload](#) section above, restart functionality is implemented by using two classloaders. For most applications this approach works well, however, sometimes it can cause classloading issues.

By default, any open project in your IDE will be loaded using the “restart” classloader, and any regular `.jar` file will be loaded using the “base” classloader. If you work on a multi-module project, and not each module is imported into your IDE, you may need to customize things. To do this you can create a `META-INF/spring-devtools.properties` file.

The `spring-devtools.properties` file can contain `restart.exclude.` and `restart.include.` prefixed properties. The `include` elements are items that should be pulled up into the “restart” classloader, and the `exclude` elements are items that should be pushed down into the “base” classloader. The value of the property is a regex pattern that will be applied to the classpath.

For example:

```
restart.exclude.companycommonlibs=/mycorp-common-[\\w-]+\\.jar
restart.include.projectcommon=/mycorp-myproj-[\\w-]+\\.jar
```

Note

All property keys must be unique. As long as a property starts with `restart.include.` or `restart.exclude.` it will be considered.

Tip

All `META-INF/spring-devtools.properties` from the classpath will be loaded. You can package files inside your project, or in the libraries that the project consumes.

Known limitations

Restart functionality does not work well with objects that are deserialized using a standard `ObjectInputStream`. If you need to deserialize data, you may need to use Spring's `ConfigurableObjectInputStream` in combination with `Thread.currentThread().getContextClassLoader()`.

Unfortunately, several third-party libraries deserialize without considering the context classloader. If you find such a problem, you will need to request a fix with the original authors.

20.3 LiveReload

The `spring-boot-devtools` module includes an embedded LiveReload server that can be used to trigger a browser refresh when a resource is changed. LiveReload browser extensions are freely available for Chrome, Firefox and Safari from livereload.com.

If you don't want to start the LiveReload server when your application runs you can set the `spring.devtools.livereload.enabled` property to `false`.

Note

You can only run one LiveReload server at a time. Before starting your application, ensure that no other LiveReload servers are running. If you start multiple applications from your IDE, only the first will have LiveReload support.

20.4 Global settings

You can configure global devtools settings by adding a file named `.spring-boot-devtools.properties` to your `$HOME` folder (note that the filename starts with `.`). Any properties added to this file will apply to *all* Spring Boot applications on your machine that use devtools. For example, to configure restart to always use a [trigger file](#), you would add the following:

`~/.spring-boot-devtools.properties`.

```
spring.devtools.reload.trigger-file=.reloadtrigger
```

20.5 Remote applications

The Spring Boot developer tools are not just limited to local development. You can also use several features when running applications remotely. Remote support is opt-in, to enable it you need to make sure that `devtools` is included in the repackaged archive:

```
<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
      <configuration>
        <excludeDevtools>>false</excludeDevtools>
      </configuration>
    </plugin>
  </plugins>
</build>
```

Then you need to set a `spring.devtools.remote.secret` property, for example:

```
spring.devtools.remote.secret=mysecret
```

Warning

Enabling `spring-boot-devtools` on a remote application is a security risk. You should never enable support on a production deployment.

Tip

If you need to use a proxy to access the remote application, configure the `spring.devtools.remote.proxy.host` and `spring.devtools.remote.proxy.port` properties.

Remote update

The remote client will monitor your application classpath for changes in the same way as the [local restart](#). Any updated resource will be pushed to the remote application and (*if required*) trigger a restart. This can be quite helpful if you are iterating on a feature that uses a cloud service that you don't have locally. Generally remote updates and restarts are much quicker than a full rebuild and deploy cycle.

Note

Files are only monitored when the remote client is running. If you change a file before starting the remote client, it won't be pushed to the remote server.

21. Packaging your application for production

Executable jars can be used for production deployment. As they are self-contained, they are also ideally suited for cloud-based deployment.

For additional “production ready” features, such as health, auditing and metric REST or JMX endpoints; consider adding `spring-boot-actuator`. See [Part V, “Spring Boot Actuator: Production-ready features”](#) for details.

22. What to read next

You should now have good understanding of how you can use Spring Boot along with some best practices that you should follow. You can now go on to learn about specific [Spring Boot features](#) in depth, or you could skip ahead and read about the “[production ready](#)” aspects of Spring Boot.

Part IV. Spring Boot features

This section dives into the details of Spring Boot. Here you can learn about the key features that you will want to use and customize. If you haven't already, you might want to read the [Part II, "Getting started"](#) and [Part III, "Using Spring Boot"](#) sections so that you have a good grounding of the basics.

so you need to [enable the debug property](#) or [enable DEBUG logging](#) for `org.springframework.boot.autoconfigure.logging.AutoConfigurationReportLoggingInitializer`.

For instance, if you are running your application using `java -jar` you can enable the `debug` property as follows:

```
$ java -jar myproject-0.0.1-SNAPSHOT.jar --debug
```

23.2 Customizing the Banner

The banner that is printed on start up can be changed by adding a `banner.txt` file to your classpath, or by setting `banner.location` to the location of such a file. If the file has an unusual encoding you can set `banner.charset` (default is `UTF-8`). In addition to a text file, you can also add a `banner.gif`, `banner.jpg` or `banner.png` image file to your classpath, or set a `banner.image.location` property. Images will be converted into an ASCII art representation and printed above any text banner.

Inside your `banner.txt` file you can use any of the following placeholders:

Table 23.1. Banner variables

Variable	Description
<code>\${application.version}</code>	The version number of your application as declared in <code>MANIFEST.MF</code> . For example <code>Implementation-Version: 1.0</code> is printed as <code>1.0</code> .
<code>\${application.formatted-version}</code>	The version number of your application as declared in <code>MANIFEST.MF</code> formatted for display (surrounded with brackets and prefixed with <code>v</code>). For example <code>(v1.0)</code> .
<code>\${spring-boot.version}</code>	The Spring Boot version that you are using. For example <code>2.0.0.M3</code> .
<code>\${spring-boot.formatted-version}</code>	The Spring Boot version that you are using formatted for display (surrounded with brackets and prefixed with <code>v</code>). For example <code>(v2.0.0.M3)</code> .
<code>\${Ansi.NAME}</code> (or <code>\${AnsiColor.NAME}</code> , <code>\${AnsiBackground.NAME}</code> , <code>\${AnsiStyle.NAME}</code>)	Where <code>NAME</code> is the name of an ANSI escape code. See AnsiPropertySource for details.
<code>\${application.title}</code>	The title of your application as declared in <code>MANIFEST.MF</code> . For example <code>Implementation-Title: MyApp</code> is printed as <code>MyApp</code> .

Tip

The `SpringApplication.setBanner(...)` method can be used if you want to generate a banner programmatically. Use the `org.springframework.boot.Banner` interface and implement your own `printBanner()` method.

You can also use the `spring.main.banner-mode` property to determine if the banner has to be printed on `System.out` (console), using the configured logger (log) or not at all (off).

The printed banner will be registered as a singleton bean under the name `springBootBanner`.

Note

YAML maps `off` to `false` so make sure to add quotes if you want to disable the banner in your application.

```
spring:
  main:
    banner-mode: "off"
```

23.3 Customizing SpringApplication

If the `SpringApplication` defaults aren't to your taste you can instead create a local instance and customize it. For example, to turn off the banner you would write:

```
public static void main(String[] args) {
    SpringApplication app = new SpringApplication(MySpringConfiguration.class);
    app.setBannerMode(Banner.Mode.OFF);
    app.run(args);
}
```

Note

The constructor arguments passed to `SpringApplication` are configuration sources for spring beans. In most cases these will be references to `@Configuration` classes, but they could also be references to XML configuration or to packages that should be scanned.

It is also possible to configure the `SpringApplication` using an `application.properties` file. See [Chapter 24, Externalized Configuration](#) for details.

For a complete list of the configuration options, see the [SpringApplication Javadoc](#).

23.4 Fluent builder API

If you need to build an `ApplicationContext` hierarchy (multiple contexts with a parent/child relationship), or if you just prefer using a 'fluent' builder API, you can use the `SpringApplicationBuilder`.

The `SpringApplicationBuilder` allows you to chain together multiple method calls, and includes parent and child methods that allow you to create a hierarchy.

For example:

```
new SpringApplicationBuilder()
    .sources(Parent.class)
    .child(Application.class)
    .bannerMode(Banner.Mode.OFF)
    .run(args);
```

Note

There are some restrictions when creating an `ApplicationContext` hierarchy, e.g. Web components **must** be contained within the child context, and the same `Environment` will be used for both parent and child contexts. See the [SpringApplicationBuilder Javadoc](#) for full details.

23.5 Application events and listeners

In addition to the usual Spring Framework events, such as [ContextRefreshedEvent](#), a `SpringApplication` sends some additional application events.

Note

Some events are actually triggered before the `ApplicationContext` is created so you cannot register a listener on those as a `@Bean`. You can register them via the `SpringApplication.addListeners(...)` or `SpringApplicationBuilder.listeners(...)` methods.

If you want those listeners to be registered automatically regardless of the way the application is created you can add a `META-INF/spring.factories` file to your project and reference your listener(s) using the `org.springframework.context.ApplicationListener` key.

```
org.springframework.context.ApplicationListener=com.example.project.MyListener
```

Application events are sent in the following order, as your application runs:

1. An `ApplicationStartingEvent` is sent at the start of a run, but before any processing except the registration of listeners and initializers.
2. An `ApplicationEnvironmentPreparedEvent` is sent when the `Environment` to be used in the context is known, but before the context is created.
3. An `ApplicationPreparedEvent` is sent just before the refresh is started, but after bean definitions have been loaded.
4. An `ApplicationReadyEvent` is sent after the refresh and any related callbacks have been processed to indicate the application is ready to service requests.
5. An `ApplicationFailedEvent` is sent if there is an exception on startup.

Tip

You often won't need to use application events, but it can be handy to know that they exist. Internally, Spring Boot uses events to handle a variety of tasks.

23.6 Web environment

A `SpringApplication` will attempt to create the right type of `ApplicationContext` on your behalf. By default, an `AnnotationConfigApplicationContext` or `AnnotationConfigServletWebServerApplicationContext` will be used, depending on whether you are developing a web application or not.

The algorithm used to determine a 'web environment' is fairly simplistic (based on the presence of a few classes). You can use `setWebEnvironment(boolean webEnvironment)` if you need to override the default.

It is also possible to take complete control of the `ApplicationContext` type that will be used by calling `setApplicationContextClass(...)`.

Tip

It is often desirable to call `setWebEnvironment(false)` when using `SpringApplication` within a JUnit test.

23.7 Accessing application arguments

If you need to access the application arguments that were passed to `SpringApplication.run(...)` you can inject a `org.springframework.boot.ApplicationArguments` bean. The `ApplicationArguments` interface provides access to both the raw `String[]` arguments as well as parsed option and non-option arguments:

```
import org.springframework.boot.*
import org.springframework.beans.factory.annotation.*
import org.springframework.stereotype.*

@Component
public class MyBean {

    @Autowired
    public MyBean(ApplicationArguments args) {
        boolean debug = args.containsOption("debug");
        List<String> files = args.getNonOptionArgs();
        // if run with "--debug logfile.txt" debug=true, files=["logfile.txt"]
    }
}
```

Tip

Spring Boot will also register a `CommandLinePropertySource` with the Spring Environment. This allows you to also inject single application arguments using the `@Value` annotation.

23.8 Using the ApplicationRunner or CommandLineRunner

If you need to run some specific code once the `SpringApplication` has started, you can implement the `ApplicationRunner` or `CommandLineRunner` interfaces. Both interfaces work in the same way and offer a single `run` method which will be called just before `SpringApplication.run(...)` completes.

The `CommandLineRunner` interfaces provides access to application arguments as a simple string array, whereas the `ApplicationRunner` uses the `ApplicationArguments` interface discussed above.

```
import org.springframework.boot.*
import org.springframework.stereotype.*

@Component
public class MyBean implements CommandLineRunner {
```

```

public void run(String... args) {
    // Do something...
}
}

```

You can additionally implement the `org.springframework.core.Ordered` interface or use the `org.springframework.core.annotation.Order` annotation if several `CommandLineRunner` or `ApplicationRunner` beans are defined that must be called in a specific order.

23.9 Application exit

Each `SpringApplication` will register a shutdown hook with the JVM to ensure that the `ApplicationContext` is closed gracefully on exit. All the standard Spring lifecycle callbacks (such as the `DisposableBean` interface, or the `@PreDestroy` annotation) can be used.

In addition, beans may implement the `org.springframework.boot.ExitCodeGenerator` interface if they wish to return a specific exit code when `SpringApplication.exit()` is called. This exit code can then be passed to `System.exit()` to return it as a status code.

```

@SpringBootApplication
public class ExitCodeApplication {

    @Bean
    public ExitCodeGenerator exitCodeGenerator() {
        return () -> 42;
    }

    public static void main(String[] args) {
        System.exit(SpringApplication
            .exit(SpringApplication.run(ExitCodeApplication.class, args)));
    }
}

```

Also, the `ExitCodeGenerator` interface may be implemented by exceptions. When such an exception is encountered, Spring Boot will return the exit code provided by the implemented `getExitCode()` method.

23.10 Admin features

It is possible to enable admin-related features for the application by specifying the `spring.application.admin.enabled` property. This exposes the [SpringApplicationAdminMXBean](#) on the platform `MBeanServer`. You could use this feature to administer your Spring Boot application remotely. This could also be useful for any service wrapper implementation.

Tip

If you want to know on which HTTP port the application is running, get the property with key `local.server.port`.

Note

Take care when enabling this feature as the `MBean` exposes a method to shutdown the application.

24. Externalized Configuration

Spring Boot allows you to externalize your configuration so you can work with the same application code in different environments. You can use properties files, YAML files, environment variables and command-line arguments to externalize configuration. Property values can be injected directly into your beans using the `@Value` annotation, accessed via Spring's `Environment` abstraction or [bound to structured objects](#) via `@ConfigurationProperties`.

Spring Boot uses a very particular `PropertySource` order that is designed to allow sensible overriding of values. Properties are considered in the following order:

1. [Devtools global settings properties](#) on your home directory (`~/ .spring-boot-devtools.properties` when devtools is active).
2. [@TestPropertySource](#) annotations on your tests.
3. [@SpringBootTest#properties](#) annotation attribute on your tests.
4. Command line arguments.
5. Properties from `SPRING_APPLICATION_JSON` (inline JSON embedded in an environment variable or system property)
6. `ServletConfig` init parameters.
7. `ServletContext` init parameters.
8. JNDI attributes from `java:comp/env`.
9. Java System properties (`System.getProperties()`).
10. OS environment variables.
11. A `RandomValuePropertySource` that only has properties in `random.*`.
12. [Profile-specific application properties](#) outside of your packaged jar (`application-{profile}.properties` and YAML variants)
13. [Profile-specific application properties](#) packaged inside your jar (`application-{profile}.properties` and YAML variants)
14. Application properties outside of your packaged jar (`application.properties` and YAML variants).
15. Application properties packaged inside your jar (`application.properties` and YAML variants).
16. [@PropertySource](#) annotations on your `@Configuration` classes.
17. Default properties (specified using `SpringApplication.setDefaultProperties`).

To provide a concrete example, suppose you develop a `@Component` that uses a name property:

```
import org.springframework.stereotype.*
import org.springframework.beans.factory.annotation.*
```

```

@Component
public class MyBean {

    @Value("${name}")
    private String name;

    // ...

}

```

On your application classpath (e.g. inside your jar) you can have an `application.properties` that provides a sensible default property value for `name`. When running in a new environment, an `application.properties` can be provided outside of your jar that overrides the `name`; and for one-off testing, you can launch with a specific command line switch (e.g. `java -jar app.jar --name="Spring"`).

Tip

The `SPRING_APPLICATION_JSON` properties can be supplied on the command line with an environment variable. For example in a UN*X shell:

```
$ SPRING_APPLICATION_JSON='{ "foo":{ "bar": "spam" } }' java -jar myapp.jar
```

In this example you will end up with `foo.bar=spam` in the Spring Environment. You can also supply the JSON as `spring.application.json` in a System variable:

```
$ java -Dspring.application.json='{ "foo": "bar" }' -jar myapp.jar
```

or command line argument:

```
$ java -jar myapp.jar --spring.application.json='{ "foo": "bar" }'
```

or as a JNDI variable `java:comp/env/spring.application.json`.

24.1 Configuring random values

The `RandomValuePropertySource` is useful for injecting random values (e.g. into secrets or test cases). It can produce integers, longs, uuids or strings, e.g.

```

my.secret=${random.value}
my.number=${random.int}
my.bignumber=${random.long}
my.uuid=${random.uuid}
my.number.less.than.ten=${random.int(10)}
my.number.in.range=${random.int[1024,65536]}

```

The `random.int*` syntax is `OPEN value (,max) CLOSE` where the `OPEN`, `CLOSE` are any character and `value`, `max` are integers. If `max` is provided then `value` is the minimum value and `max` is the maximum (exclusive).

24.2 Accessing command line properties

By default `SpringApplication` will convert any command line option arguments (starting with `--`, e.g. `--server.port=9000`) to a property and add it to the Spring Environment. As mentioned above, command line properties always take precedence over other property sources.

If you don't want command line properties to be added to the `Environment` you can disable them using `SpringApplication.setAddCommandLineProperties(false)`.

24.3 Application property files

`SpringApplication` will load properties from `application.properties` files in the following locations and add them to the Spring `Environment`:

1. A `/config` subdirectory of the current directory.
2. The current directory
3. A classpath `/config` package
4. The classpath root

The list is ordered by precedence (properties defined in locations higher in the list override those defined in lower locations).

Note

You can also [use YAML \('.yml'\) files](#) as an alternative to '.properties'.

If you don't like `application.properties` as the configuration file name you can switch to another by specifying a `spring.config.name` environment property. You can also refer to an explicit location using the `spring.config.location` environment property (comma-separated list of directory locations, or file paths).

```
$ java -jar myproject.jar --spring.config.name=myproject
```

or

```
$ java -jar myproject.jar --spring.config.location=classpath://default.properties,classpath://override.properties
```

Warning

`spring.config.name` and `spring.config.location` are used very early to determine which files have to be loaded so they have to be defined as an environment property (typically OS env, system property or command line argument).

If `spring.config.location` contains directories (as opposed to files) they should end in `/` (and will be appended with the names generated from `spring.config.name` before being loaded, including profile-specific file names). Files specified in `spring.config.location` are used as-is, with no support for profile-specific variants, and will be overridden by any profile-specific properties.

Config locations are searched in reverse order. By default, the configured locations are `classpath://,classpath://config/,file:./,file:./config/`. The resulting search order is:

1. `file:./config/`
2. `file:./`

3. `classpath:/config/`

4. `classpath:/`

When custom config locations are configured, they are used in addition to the default locations. Custom locations are searched before the default locations. For example, if custom locations `classpath:/custom-config/`, `file:./custom-config/` are configured, the search order becomes:

1. `file:./custom-config/`

2. `classpath:custom-config/`

3. `file:/config/`

4. `file:/`

5. `classpath:/config/`

6. `classpath:/`

This search ordering allows you to specify default values in one configuration file and then selectively override those values in another. You can provide default values for your application in `application.properties` (or whatever other basename you choose with `spring.config.name`) in one of the default locations. These default values can then be overridden at runtime with a different file located in one of the custom locations.

Note

If you use environment variables rather than system properties, most operating systems disallow period-separated key names, but you can use underscores instead (e.g. `SPRING_CONFIG_NAME` instead of `spring.config.name`).

Note

If you are running in a container then JNDI properties (in `java:comp/env`) or servlet context initialization parameters can be used instead of, or as well as, environment variables or system properties.

24.4 Profile-specific properties

In addition to `application.properties` files, profile-specific properties can also be defined using the naming convention `application-{profile}.properties`. The Environment has a set of default profiles (by default `[default]`) which are used if no active profiles are set (i.e. if no profiles are explicitly activated then properties from `application-default.properties` are loaded).

Profile-specific properties are loaded from the same locations as standard `application.properties`, with profile-specific files always overriding the non-specific ones irrespective of whether the profile-specific files are inside or outside your packaged jar.

If several profiles are specified, a last wins strategy applies. For example, profiles specified by the `spring.profiles.active` property are added after those configured via the `SpringApplication` API and therefore take precedence.

Note

If you have specified any files in `spring.config.location`, profile-specific variants of those files will not be considered. Use directories in `spring.config.location` if you also want to also use profile-specific properties.

24.5 Placeholders in properties

The values in `application.properties` are filtered through the existing `Environment` when they are used so you can refer back to previously defined values (e.g. from System properties).

```
app.name=MyApp
app.description=${app.name} is a Spring Boot application
```

Tip

You can also use this technique to create ‘short’ variants of existing Spring Boot properties. See the [Section 73.4, “Use ‘short’ command line arguments”](#) how-to for details.

24.6 Using YAML instead of Properties

[YAML](#) is a superset of JSON, and as such is a very convenient format for specifying hierarchical configuration data. The `SpringApplication` class will automatically support YAML as an alternative to properties whenever you have the [SnakeYAML](#) library on your classpath.

Note

If you use ‘Starters’ SnakeYAML will be automatically provided via `spring-boot-starter`.

Loading YAML

Spring Framework provides two convenient classes that can be used to load YAML documents. The `YamlPropertiesFactoryBean` will load YAML as `Properties` and the `YamlMapFactoryBean` will load YAML as a `Map`.

For example, the following YAML document:

```
environments:
  dev:
    url: http://dev.bar.com
    name: Developer Setup
  prod:
    url: http://foo.bar.com
    name: My Cool App
```

Would be transformed into these properties:

```
environments.dev.url=http://dev.bar.com
environments.dev.name=Developer Setup
environments.prod.url=http://foo.bar.com
environments.prod.name=My Cool App
```

YAML lists are represented as property keys with `[index]` dereferencers, for example this YAML:

```
my:
  servers:
```

```
- dev.bar.com
- foo.bar.com
```

Would be transformed into these properties:

```
my.servers[0]=dev.bar.com
my.servers[1]=foo.bar.com
```

To bind to properties like that using the `Spring DataBinder` utilities (which is what `@ConfigurationProperties` does) you need to have a property in the target bean of type `java.util.List` (or `Set`) and you either need to provide a setter, or initialize it with a mutable value, e.g. this will bind to the properties above

```
@ConfigurationProperties(prefix="my")
public class Config {

    private List<String> servers = new ArrayList<String>();

    public List<String> getServers() {
        return this.servers;
    }
}
```

Note

Extra care is required when configuring lists that way as overriding will not work as you would expect. In the example above, when `my.servers` is redefined in several places, the individual elements are targeted for override, not the list. To make sure that a `PropertySource` with higher precedence can override the list, you need to define it as a single property:

```
my:
  servers: dev.bar.com,foo.bar.com
```

Exposing YAML as properties in the Spring Environment

The `YamlPropertySourceLoader` class can be used to expose YAML as a `PropertySource` in the Spring Environment. This allows you to use the familiar `@Value` annotation with placeholders syntax to access YAML properties.

Multi-profile YAML documents

You can specify multiple profile-specific YAML documents in a single file by using a `spring.profiles` key to indicate when the document applies. For example:

```
server:
  address: 192.168.1.100
---
spring:
  profiles: development
server:
  address: 127.0.0.1
---
spring:
  profiles: production
server:
  address: 192.168.1.120
```

In the example above, the `server.address` property will be `127.0.0.1` if the `development` profile is active. If the `development` and `production` profiles are **not** enabled, then the value for the property will be `192.168.1.100`.

The default profiles are activated if none are explicitly active when the application context starts. So in this YAML we set a value for `security.user.password` that is **only** available in the "default" profile:

```
server:
  port: 8000
---
spring:
  profiles: default
security:
  user:
    password: weak
```

whereas in this example, the password is always set because it isn't attached to any profile, and it would have to be explicitly reset in all other profiles as necessary:

```
server:
  port: 8000
security:
  user:
    password: weak
```

Spring profiles designated using the "spring.profiles" element may optionally be negated using the `!` character. If both negated and non-negated profiles are specified for a single document, at least one non-negated profile must match and no negated profiles may match.

YAML shortcomings

YAML files can't be loaded via the `@PropertySource` annotation. So in the case that you need to load values that way, you need to use a properties file.

Merging YAML lists

As [we have seen above](#), any YAML content is ultimately transformed to properties. That process may be counter intuitive when overriding "list" properties via a profile.

For example, assume a `MyPojo` object with `name` and `description` attributes that are `null` by default. Let's expose a list of `MyPojo` from `FooProperties`:

```
@ConfigurationProperties("foo")
public class FooProperties {

    private final List<MyPojo> list = new ArrayList<>();

    public List<MyPojo> getList() {
        return this.list;
    }

}
```

Consider the following configuration:

```
foo:
  list:
    - name: my name
      description: my description
---
spring:
  profiles: dev
foo:
  list:
    - name: my another name
```

If the `dev` profile isn't active, `FooProperties.list` will contain one `MyPojo` entry as defined above. If the `dev` profile is enabled however, the `list` will *still* only contain one entry (with name "my another name" and description `null`). This configuration *will not* add a second `MyPojo` instance to the list, and it won't merge the items.

When a collection is specified in multiple profiles, the one with highest priority is used (and only that one):

```
foo:
  list:
    - name: my name
      description: my description
    - name: another name
      description: another description
---
spring:
  profiles: dev
foo:
  list:
    - name: my another name
```

In the example above, considering that the `dev` profile is active, `FooProperties.list` will contain *one* `MyPojo` entry (with name "my another name" and description `null`).

24.7 Type-safe Configuration Properties

Using the `@Value("${property}")` annotation to inject configuration properties can sometimes be cumbersome, especially if you are working with multiple properties or your data is hierarchical in nature. Spring Boot provides an alternative method of working with properties that allows strongly typed beans to govern and validate the configuration of your application.

```
package com.example;

import java.net.InetAddress;
import java.util.ArrayList;
import java.util.Collections;
import java.util.List;

import org.springframework.boot.context.properties.ConfigurationProperties;

@ConfigurationProperties("foo")
public class FooProperties {

    private boolean enabled;

    private InetAddress remoteAddress;

    private final Security security = new Security();

    public boolean isEnabled() { ... }

    public void setEnabled(boolean enabled) { ... }

    public InetAddress getRemoteAddress() { ... }

    public void setRemoteAddress(InetAddress remoteAddress) { ... }

    public Security getSecurity() { ... }

    public static class Security {

        private String username;

        private String password;

        private List<String> roles = new ArrayList<>(Collections.singleton("USER"));
```

```

    public String getUsername() { ... }

    public void setUsername(String username) { ... }

    public String getPassword() { ... }

    public void setPassword(String password) { ... }

    public List<String> getRoles() { ... }

    public void setRoles(List<String> roles) { ... }
}

```

The POJO above defines the following properties:

- `foo.enabled`, false by default
- `foo.remote-address`, with a type that can be coerced from `String`
- `foo.security.username`, with a nested "security" whose name is determined by the name of the property. In particular the return type is not used at all there and could have been `SecurityProperties`
- `foo.security.password`
- `foo.security.roles`, with a collection of `String`

Note

Getters and setters are usually mandatory, since binding is via standard Java Beans property descriptors, just like in Spring MVC. There are cases where a setter may be omitted:

- Maps, as long as they are initialized, need a getter but not necessarily a setter since they can be mutated by the binder.
- Collections and arrays can be either accessed via an index (typically with YAML) or using a single comma-separated value (properties). In the latter case, a setter is mandatory. We recommend to always add a setter for such types. If you initialize a collection, make sure it is not immutable (as in the example above)
- If nested POJO properties are initialized (like the `security` field in the example above), a setter is not required. If you want the binder to create the instance on-the-fly using its default constructor, you will need a setter.

Some people use Project Lombok to add getters and setters automatically. Make sure that Lombok doesn't generate any particular constructor for such type as it will be used automatically by the container to instantiate the object.

Tip

See also the [differences between `@Value` and `@ConfigurationProperties`](#).

You also need to list the properties classes to register in the `@EnableConfigurationProperties` annotation:

```
@Configuration
```

```
@EnableConfigurationProperties(FooProperties.class)
public class MyConfiguration {
}
```

Note

When `@ConfigurationProperties` bean is registered that way, the bean will have a conventional name: `<prefix>-<fqcn>`, where `<prefix>` is the environment key prefix specified in the `@ConfigurationProperties` annotation and `<fqcn>` the fully qualified name of the bean. If the annotation does not provide any prefix, only the fully qualified name of the bean is used.

The bean name in the example above will be `foo-com.example.FooProperties`.

Even if the configuration above will create a regular bean for `FooProperties`, we recommend that `@ConfigurationProperties` only deal with the environment and in particular does not inject other beans from the context. Having said that, The `@EnableConfigurationProperties` annotation is *also* automatically applied to your project so that any *existing* bean annotated with `@ConfigurationProperties` will be configured from the Environment. You could shortcut `MyConfiguration` above by making sure `FooProperties` is already a bean:

```
@Component
@ConfigurationProperties(prefix="foo")
public class FooProperties {

    // ... see above

}
```

This style of configuration works particularly well with the `SpringApplication` external YAML configuration:

```
# application.yml

foo:
  remote-address: 192.168.1.1
  security:
    username: foo
    roles:
      - USER
      - ADMIN

# additional configuration as required
```

To work with `@ConfigurationProperties` beans you can just inject them in the same way as any other bean.

```
@Service
public class MyService {

    private final FooProperties properties;

    @Autowired
    public MyService(FooProperties properties) {
        this.properties = properties;
    }

    //...

    @PostConstruct
    public void openConnection() {
        Server server = new Server(this.properties.getRemoteAddress());
        // ...
    }
}
```

```

    }
}

```

Tip

Using `@ConfigurationProperties` also allows you to generate meta-data files that can be used by IDEs to offer auto-completion for your own keys, see the [Appendix B, Configuration meta-data](#) appendix for details.

Third-party configuration

As well as using `@ConfigurationProperties` to annotate a class, you can also use it on public `@Bean` methods. This can be particularly useful when you want to bind properties to third-party components that are outside of your control.

To configure a bean from the `Environment` properties, add `@ConfigurationProperties` to its bean registration:

```

@ConfigurationProperties(prefix = "bar")
@Bean
public BarComponent barComponent() {
    ...
}

```

Any property defined with the `bar` prefix will be mapped onto that `BarComponent` bean in a similar manner as the `FooProperties` example above.

Relaxed binding

Spring Boot uses some relaxed rules for binding `Environment` properties to `@ConfigurationProperties` beans, so there doesn't need to be an exact match between the `Environment` property name and the bean property name. Common examples where this is useful include dashed separated (e.g. `context-path` binds to `contextPath`), and capitalized (e.g. `PORT` binds to `port`) environment properties.

For example, given the following `@ConfigurationProperties` class:

```

@ConfigurationProperties(prefix="person")
public class OwnerProperties {

    private String firstName;

    public String getFirstName() {
        return this.firstName;
    }

    public void setFirstName(String firstName) {
        this.firstName = firstName;
    }

}

```

The following properties names can all be used:

Table 24.1. relaxed binding

Property	Note
<code>person.firstName</code>	Standard camel case syntax.

Property	Note
<code>person.first-name</code>	Kebab-case, recommended for use in <code>.properties</code> and <code>.yml</code> files.
<code>person.first_name</code>	Underscore notation, alternative format for use in <code>.properties</code> and <code>.yml</code> files.
<code>PERSON_FIRSTNAME</code>	Upper case format. Recommended when using a system environment variables.

Note

The `prefix` value for the annotation must be in kebab-case, ie, lowercase and separated by `-`.

Table 24.2. relaxed binding rules per property source

Property Source	Simple	List
Properties Files	Camel-case, kebab-case or underscore notation	Standard list syntax using <code>[]</code> or comma-separated values
YAML Files	Camel-case, kebab-case or underscore notation	Standard YAML list syntax or comma-separated values
Environment Variables	Upper case format with underscore as the delimiter. <code>_</code> should not be used within a property name	Numeric values surrounded by underscores. eg: <code>MY_FOO_1_BAR = my.foo[1].bar</code>
System properties	Camel-case, kebab-case or underscore notation	Standard list syntax using <code>[]</code> or comma-separated values

Tip

We recommend that, when possible, properties are stored in lowercase kebab format. i.e. `my.property-name=foo`

Properties conversion

Spring will attempt to coerce the external application properties to the right type when it binds to the `@ConfigurationProperties` beans. If you need custom type conversion you can provide a `ConversionService` bean (with bean id `conversionService`) or custom property editors (via a `CustomEditorConfigurer` bean) or custom `Converters` (with bean definitions annotated as `@ConfigurationPropertiesBinding`).

Note

As this bean is requested very early during the application lifecycle, make sure to limit the dependencies that your `ConversionService` is using. Typically, any dependency that you require may not be fully initialized at creation time. You may want to rename your custom `ConversionService` if it's not required for configuration keys coercion and only rely on custom converters qualified with `@ConfigurationPropertiesBinding`.

@ConfigurationProperties Validation

Spring Boot will attempt to validate `@ConfigurationProperties` classes whenever they are annotated with Spring's `@Validated` annotation. You can use JSR-303 `javax.validation` constraint annotations directly on your configuration class. Simply ensure that a compliant JSR-303 implementation is on your classpath, then add constraint annotations to your fields:

```
@ConfigurationProperties(prefix="foo")
@Validated
public class FooProperties {

    @NotNull
    private InetAddress remoteAddress;

    // ... getters and setters

}
```

In order to validate values of nested properties, you must annotate the associated field as `@Valid` to trigger its validation. For example, building upon the above `FooProperties` example:

```
@ConfigurationProperties(prefix="connection")
@Validated
public class FooProperties {

    @NotNull
    private InetAddress remoteAddress;

    @Valid
    private final Security security = new Security();

    // ... getters and setters

    public static class Security {

        @NotEmpty
        public String username;

        // ... getters and setters

    }

}
```

You can also add a custom Spring Validator by creating a bean definition called `configurationPropertiesValidator`. The `@Bean` method should be declared `static`. The configuration properties validator is created very early in the application's lifecycle and declaring the `@Bean` method as `static` allows the bean to be created without having to instantiate the `@Configuration` class. This avoids any problems that may be caused by early instantiation. There is a [property validation sample](#) so you can see how to set things up.

Tip

The `spring-boot-actuator` module includes an endpoint that exposes all `@ConfigurationProperties` beans. Simply point your web browser to `/configprops` or use the equivalent JMX endpoint. See the [Production ready features](#) section for details.

@ConfigurationProperties vs. @Value

@Value is a core container feature and it does not provide the same features as type-safe Configuration Properties. The table below summarizes the features that are supported by @ConfigurationProperties and @Value:

Feature	@ConfigurationProperties	@Value
Relaxed binding	Yes	No
Meta-data support	Yes	No
SpEL evaluation	No	Yes

If you define a set of configuration keys for your own components, we recommend you to group them in a POJO annotated with @ConfigurationProperties. Please also be aware that since @Value does not support relaxed binding, it isn't a great candidate if you need to provide the value using environment variables.

Finally, while you can write a SpEL expression in @Value, such expressions are not processed from [Application property files](#).

25. Profiles

Spring Profiles provide a way to segregate parts of your application configuration and make it only available in certain environments. Any `@Component` or `@Configuration` can be marked with `@Profile` to limit when it is loaded:

```
@Configuration
@Profile("production")
public class ProductionConfiguration {

    // ...

}
```

In the normal Spring way, you can use a `spring.profiles.active` Environment property to specify which profiles are active. You can specify the property in any of the usual ways, for example you could include it in your `application.properties`:

```
spring.profiles.active=dev,hsqldb
```

or specify on the command line using the switch `--spring.profiles.active=dev,hsqldb`.

25.1 Adding active profiles

The `spring.profiles.active` property follows the same ordering rules as other properties, the highest `PropertySource` will win. This means that you can specify active profiles in `application.properties` then **replace** them using the command line switch.

Sometimes it is useful to have profile-specific properties that **add** to the active profiles rather than replace them. The `spring.profiles.include` property can be used to unconditionally add active profiles. The `SpringApplication` entry point also has a Java API for setting additional profiles (i.e. on top of those activated by the `spring.profiles.active` property): see the `setAdditionalProfiles()` method.

For example, when an application with following properties is run using the switch `--spring.profiles.active=prod` the `proddb` and `prodmq` profiles will also be activated:

```
---
my.property: fromyamlfile
---
spring.profiles: prod
spring.profiles.include:
- proddb
- prodmq
```

Note

Remember that the `spring.profiles` property can be defined in a YAML document to determine when this particular document is included in the configuration. See [Section 73.7](#), “Change configuration depending on the environment” for more details.

25.2 Programmatically setting profiles

You can programmatically set active profiles by calling `SpringApplication.setAdditionalProfiles(...)` before your application runs. It is also possible to activate profiles using Spring’s `ConfigurableEnvironment` interface.

25.3 Profile-specific configuration files

Profile-specific variants of both `application.properties` (or `application.yml`) and files referenced via `@ConfigurationProperties` are considered as files are loaded. See [Section 24.4, “Profile-specific properties”](#) for details.

26. Logging

Spring Boot uses [Commons Logging](#) for all internal logging, but leaves the underlying log implementation open. Default configurations are provided for [Java Util Logging](#), [Log4J2](#) and [Logback](#). In each case loggers are pre-configured to use console output with optional file output also available.

By default, if you use the ‘Starters’, Logback will be used for logging. Appropriate Logback routing is also included to ensure that dependent libraries that use Java Util Logging, Commons Logging, Log4J or SLF4J will all work correctly.

Tip

There are a lot of logging frameworks available for Java. Don't worry if the above list seems confusing. Generally you won't need to change your logging dependencies and the Spring Boot defaults will work just fine.

26.1 Log format

The default log output from Spring Boot looks like this:

```
2014-03-05 10:57:51.112 INFO 45469 --- [          main] org.apache.catalina.core.StandardEngine :
Starting Servlet Engine: Apache Tomcat/7.0.52
2014-03-05 10:57:51.253 INFO 45469 --- [ost-startStop-1] o.a.c.c.C.[Tomcat].[localhost].[/] :
Initializing Spring embedded WebApplicationContext
2014-03-05 10:57:51.253 INFO 45469 --- [ost-startStop-1] o.s.web.context.ContextLoader :
Root WebApplicationContext: initialization completed in 1358 ms
2014-03-05 10:57:51.698 INFO 45469 --- [ost-startStop-1] o.s.b.c.e.ServletRegistrationBean :
Mapping servlet: 'dispatcherServlet' to [/]
2014-03-05 10:57:51.702 INFO 45469 --- [ost-startStop-1] o.s.b.c.embedded.FilterRegistrationBean :
Mapping filter: 'hiddenHttpMethodFilter' to: [/*]
```

The following items are output:

- Date and Time — Millisecond precision and easily sortable.
- Log Level — ERROR, WARN, INFO, DEBUG or TRACE.
- Process ID.
- A --- separator to distinguish the start of actual log messages.
- Thread name — Enclosed in square brackets (may be truncated for console output).
- Logger name — This is usually the source class name (often abbreviated).
- The log message.

Note

Logback does not have a FATAL level (it is mapped to ERROR)

26.2 Console output

The default log configuration will echo messages to the console as they are written. By default ERROR, WARN and INFO level messages are logged. You can also enable a “debug” mode by starting your application with a `--debug` flag.

```
$ java -jar myapp.jar --debug
```

Note

you can also specify `debug=true` in your `application.properties`.

When the debug mode is enabled, a selection of core loggers (embedded container, Hibernate and Spring Boot) are configured to output more information. Enabling the debug mode does *not* configure your application to log all messages with `DEBUG` level.

Alternatively, you can enable a “trace” mode by starting your application with a `--trace` flag (or `trace=true` in your `application.properties`). This will enable trace logging for a selection of core loggers (embedded container, Hibernate schema generation and the whole Spring portfolio).

Color-coded output

If your terminal supports ANSI, color output will be used to aid readability. You can set `spring.output.ansi.enabled` to a [supported value](#) to override the auto detection.

Color coding is configured using the `%clr` conversion word. In its simplest form the converter will color the output according to the log level, for example:

```
%clr(%5p)
```

The mapping of log level to a color is as follows:

Level	Color
FATAL	Red
ERROR	Red
WARN	Yellow
INFO	Green
DEBUG	Green
TRACE	Green

Alternatively, you can specify the color or style that should be used by providing it as an option to the conversion. For example, to make the text yellow:

```
%clr(%d{yyyy-MM-dd HH:mm:ss.SSS}){yellow}
```

The following colors and styles are supported:

- blue
- cyan
- faint
- green
- magenta

- red
- yellow

26.3 File output

By default, Spring Boot will only log to the console and will not write log files. If you want to write log files in addition to the console output you need to set a `logging.file` or `logging.path` property (for example in your `application.properties`).

The following table shows how the `logging.*` properties can be used together:

Table 26.1. Logging properties

<code>logging.file</code>	<code>logging.path</code>	Example	Description
<i>(none)</i>	<i>(none)</i>		Console only logging.
Specific file	<i>(none)</i>	<code>my.log</code>	Writes to the specified log file. Names can be an exact location or relative to the current directory.
<i>(none)</i>	Specific directory	<code>/var/log</code>	Writes <code>spring.log</code> to the specified directory. Names can be an exact location or relative to the current directory.

Log files will rotate when they reach 10 MB and as with console output, `ERROR`, `WARN` and `INFO` level messages are logged by default.

Note

The logging system is initialized early in the application lifecycle and as such logging properties will not be found in property files loaded via `@PropertySource` annotations.

Tip

Logging properties are independent of the actual logging infrastructure. As a result, specific configuration keys (such as `logback.configurationFile` for Logback) are not managed by Spring Boot.

26.4 Log Levels

All the supported logging systems can have the logger levels set in the Spring Environment (so for example in `application.properties`) using `logging.level.*=LEVEL` where 'LEVEL' is one of `TRACE`, `DEBUG`, `INFO`, `WARN`, `ERROR`, `FATAL`, `OFF`. The `root` logger can be configured using `logging.level.root`. Example `application.properties`:

```
logging.level.root=WARN
logging.level.org.springframework.web=DEBUG
logging.level.org.hibernate=ERROR
```

Note

By default Spring Boot remaps Thymeleaf `INFO` messages so that they are logged at `DEBUG` level. This helps to reduce noise in the standard log output. See [LevelRemappingAppender](#) for details of how you can apply remapping in your own configuration.

26.5 Custom log configuration

The various logging systems can be activated by including the appropriate libraries on the classpath, and further customized by providing a suitable configuration file in the root of the classpath, or in a location specified by the Spring `Environment` property `logging.config`.

You can force Spring Boot to use a particular logging system using the `org.springframework.boot.logging.LoggingSystem` system property. The value should be the fully-qualified class name of a `LoggingSystem` implementation. You can also disable Spring Boot's logging configuration entirely by using a value of `none`.

Note

Since logging is initialized **before** the `ApplicationContext` is created, it isn't possible to control logging from `@PropertySources` in Spring `@Configuration` files. System properties and the conventional Spring Boot external configuration files work just fine.)

Depending on your logging system, the following files will be loaded:

Logging System	Customization
Logback	<code>logback-spring.xml</code> , <code>logback-spring.groovy</code> , <code>logback.xml</code> or <code>logback.groovy</code>
Log4j2	<code>log4j2-spring.xml</code> or <code>log4j2.xml</code>
JDK (Java Util Logging)	<code>logging.properties</code>

Note

When possible we recommend that you use the `-spring` variants for your logging configuration (for example `logback-spring.xml` rather than `logback.xml`). If you use standard configuration locations, Spring cannot completely control log initialization.

Warning

There are known classloading issues with Java Util Logging that cause problems when running from an 'executable jar'. We recommend that you avoid it if at all possible.

To help with the customization some other properties are transferred from the Spring `Environment` to System properties:

Spring Environment	System Property	Comments
<code>logging.exception-conversion-word</code>	<code>LOG_EXCEPTION_CONVERSION_WORD</code>	The conversion word that's used when logging exceptions.
<code>logging.file</code>	<code>LOG_FILE</code>	Used in default log configuration if defined.

Spring Environment	System Property	Comments
<code>logging.path</code>	<code>LOG_PATH</code>	Used in default log configuration if defined.
<code>logging.pattern.console</code>	<code>CONSOLE_LOG_PATTERN</code>	The log pattern to use on the console (stdout). (Only supported with the default logback setup.)
<code>logging.pattern.file</code>	<code>FILE_LOG_PATTERN</code>	The log pattern to use in a file (if <code>LOG_FILE</code> enabled). (Only supported with the default logback setup.)
<code>logging.pattern.level</code>	<code>LOG_LEVEL_PATTERN</code>	The format to use to render the log level (default <code>%5p</code>). (Only supported with the default logback setup.)
<code>PID</code>	<code>PID</code>	The current process ID (discovered if possible and when not already defined as an OS environment variable).

All the logging systems supported can consult System properties when parsing their configuration files. See the default configurations in `spring-boot.jar` for examples.

Tip

If you want to use a placeholder in a logging property, you should use [Spring Boot's syntax](#) and not the syntax of the underlying framework. Notably, if you're using Logback, you should use `:` as the delimiter between a property name and its default value and not `:-`.

Tip

You can add MDC and other ad-hoc content to log lines by overriding only the `LOG_LEVEL_PATTERN` (or `logging.pattern.level` with Logback). For example, if you use `logging.pattern.level=user:%X{user} %5p` then the default log format will contain an MDC entry for "user" if it exists, e.g.

```
2015-09-30 12:30:04.031 user:juergen INFO 22174 --- [nio-8080-exec-0] demo.Controller
Handling authenticated request
```

26.6 Logback extensions

Spring Boot includes a number of extensions to Logback which can help with advanced configuration. You can use these extensions in your `logback-spring.xml` configuration file.

Note

You cannot use extensions in the standard `logback.xml` configuration file since it's loaded too early. You need to either use `logback-spring.xml` or define a `logging.config` property.

Warning

The extensions cannot be used with Logback's [configuration scanning](#). If you attempt to do so, making changes to the configuration file will result in an error similar to one of the following being logged:

```
ERROR in ch.qos.logback.core.joran.spi.Interpreter@4:71 - no applicable action for [springProperty],
current ElementPath is [[configuration][springProperty]]
ERROR in ch.qos.logback.core.joran.spi.Interpreter@4:71 - no applicable action for [springProfile],
current ElementPath is [[configuration][springProfile]]
```

Profile-specific configuration

The `<springProfile>` tag allows you to optionally include or exclude sections of configuration based on the active Spring profiles. Profile sections are supported anywhere within the `<configuration>` element. Use the `name` attribute to specify which profile accepts the configuration. Multiple profiles can be specified using a comma-separated list.

```
<springProfile name="staging">
  <!-- configuration to be enabled when the "staging" profile is active -->
</springProfile>

<springProfile name="dev, staging">
  <!-- configuration to be enabled when the "dev" or "staging" profiles are active -->
</springProfile>

<springProfile name="!production">
  <!-- configuration to be enabled when the "production" profile is not active -->
</springProfile>
```

Environment properties

The `<springProperty>` tag allows you to surface properties from the Spring Environment for use within Logback. This can be useful if you want to access values from your `application.properties` file in your logback configuration. The tag works in a similar way to Logback's standard `<property>` tag, but rather than specifying a direct `value` you specify the `source` of the property (from the Environment). You can use the `scope` attribute if you need to store the property somewhere other than in `local` scope. If you need a fallback value in case the property is not set in the Environment, you can use the `defaultValue` attribute.

```
<springProperty scope="context" name="fluentHost" source="myapp.fluentd.host"
  defaultValue="localhost"/>
<appender name="FLUENT" class="ch.qos.logback.more.appenders.DataFluentAppender">
  <remoteHost>${fluentHost}</remoteHost>
  ...
</appender>
```

Note

The `source` must be specified using kebab-case (`my.property-name`). However, properties can be added to the Environment using the relaxed rules.

27. Developing web applications

Spring Boot is well suited for web application development. You can easily create a self-contained HTTP server using embedded Tomcat, Jetty, Undertow, or Netty. Most web applications will use the `spring-boot-starter-web` module to get up and running quickly. You can also choose to use to build reactive web applications by using the `spring-boot-starter-webflux` module.

If you haven't yet developed a Spring Boot web application you can follow the "Hello World!" example in the [Getting started](#) section.

27.1 The 'Spring Web MVC framework'

The Spring Web MVC framework (often referred to as simply 'Spring MVC') is a rich 'model view controller' web framework. Spring MVC lets you create special `@Controller` or `@RestController` beans to handle incoming HTTP requests. Methods in your controller are mapped to HTTP using `@RequestMapping` annotations.

Here is a typical example `@RestController` to serve JSON data:

```
@RestController
@RequestMapping(value="/users")
public class MyRestController {

    @RequestMapping(value="/{user}", method=RequestMethod.GET)
    public User getUser(@PathVariable Long user) {
        // ...
    }

    @RequestMapping(value="/{user}/customers", method=RequestMethod.GET)
    List<Customer> getUserCustomers(@PathVariable Long user) {
        // ...
    }

    @RequestMapping(value="/{user}", method=RequestMethod.DELETE)
    public User deleteUser(@PathVariable Long user) {
        // ...
    }
}
```

Spring MVC is part of the core Spring Framework and detailed information is available in the [reference documentation](#). There are also several guides available at spring.io/guides that cover Spring MVC.

Spring MVC auto-configuration

Spring Boot provides auto-configuration for Spring MVC that works well with most applications.

The auto-configuration adds the following features on top of Spring's defaults:

- Inclusion of `ContentNegotiatingViewResolver` and `BeanNameViewResolver` beans.
- Support for serving static resources, including support for WebJars (see below).
- Automatic registration of `Converter`, `GenericConverter`, `Formatter` beans.
- Support for `HttpMessageConverters` (see below).
- Automatic registration of `MessageCodesResolver` (see below).

- Static `index.html` support.
- Custom `Favicon` support (see below).
- Automatic use of a `ConfigurableWebBindingInitializer` bean (see below).

If you want to keep Spring Boot MVC features, and you just want to add additional [MVC configuration](#) (interceptors, formatters, view controllers etc.) you can add your own `@Configuration` class of type `WebMvcConfigurer`, but **without** `@EnableWebMvc`. If you wish to provide custom instances of `RequestMappingHandlerMapping`, `RequestMappingHandlerAdapter` or `ExceptionHandlerResolver` you can declare a `WebMvcRegistrationsAdapter` instance providing such components.

If you want to take complete control of Spring MVC, you can add your own `@Configuration` annotated with `@EnableWebMvc`.

HttpMessageConverters

Spring MVC uses the `HttpMessageConverter` interface to convert HTTP requests and responses. Sensible defaults are included out of the box, for example Objects can be automatically converted to JSON (using the Jackson library) or XML (using the Jackson XML extension if available, else using JAXB). Strings are encoded using UTF-8 by default.

If you need to add or customize converters you can use Spring Boot's `HttpMessageConverters` class:

```
import org.springframework.boot.autoconfigure.web.HttpMessageConverters;
import org.springframework.context.annotation.*;
import org.springframework.http.converter.*;

@Configuration
public class MyConfiguration {

    @Bean
    public HttpMessageConverters customConverters() {
        HttpMessageConverter<?> additional = ...
        HttpMessageConverter<?> another = ...
        return new HttpMessageConverters(additional, another);
    }
}
```

Any `HttpMessageConverter` bean that is present in the context will be added to the list of converters. You can also override default converters that way.

Custom JSON Serializers and Deserializers

If you're using Jackson to serialize and deserialize JSON data, you might want to write your own `JsonSerializer` and `JsonDeserializer` classes. Custom serializers are usually [registered with Jackson via a Module](#), but Spring Boot provides an alternative `@JsonComponent` annotation which makes it easier to directly register Spring Beans.

You can use `@JsonComponent` directly on `JsonSerializer` or `JsonDeserializer` implementations. You can also use it on classes that contains serializers/deserializers as inner-classes. For example:

```
import java.io.*;
import com.fasterxml.jackson.core.*;
```

```

import com.fasterxml.jackson.databind.*;
import org.springframework.boot.jackson.*;

@JsonComponent
public class Example {

    public static class Serializer extends JsonSerializer<SomeObject> {
        // ...
    }

    public static class Deserializer extends JsonDeserializer<SomeObject> {
        // ...
    }
}

```

All `@JsonComponent` beans in the `ApplicationContext` will be automatically registered with Jackson, and since `@JsonComponent` is meta-annotated with `@Component`, the usual component-scanning rules apply.

Spring Boot also provides [JsonObjectSerializer](#) and [JsonObjectDeserializer](#) base classes which provide useful alternatives to the standard Jackson versions when serializing Objects. See the Javadoc for details.

MessageCodesResolver

Spring MVC has a strategy for generating error codes for rendering error messages from binding errors: `MessageCodesResolver`. Spring Boot will create one for you if you set the `spring.mvc.message-codes-resolver.format` property `PREFIX_ERROR_CODE` or `POSTFIX_ERROR_CODE` (see the enumeration in `DefaultMessageCodesResolver.Format`).

Static Content

By default Spring Boot will serve static content from a directory called `/static` (or `/public` or `/resources` or `/META-INF/resources`) in the classpath or from the root of the `ServletContext`. It uses the `ResourceHttpRequestHandler` from Spring MVC so you can modify that behavior by adding your own `WebMvcConfigurer` and overriding the `addResourceHandlers` method.

In a stand-alone web application the default servlet from the container is also enabled, and acts as a fallback, serving content from the root of the `ServletContext` if Spring decides not to handle it. Most of the time this will not happen (unless you modify the default MVC configuration) because Spring will always be able to handle requests through the `DispatcherServlet`.

By default, resources are mapped on `/**` but you can tune that via `spring.mvc.static-path-pattern`. For instance, relocating all resources to `/resources/**` can be achieved as follows:

```
spring.mvc.static-path-pattern=/resources/**
```

You can also customize the static resource locations using `spring.resources.static-locations` (replacing the default values with a list of directory locations). If you do this the default welcome page detection will switch to your custom locations, so if there is an `index.html` in any of your locations on startup, it will be the home page of the application.

In addition to the 'standard' static resource locations above, a special case is made for [Webjars content](#). Any resources with a path in `/webjars/**` will be served from jar files if they are packaged in the Webjars format.

Tip

Do not use the `src/main/webapp` directory if your application will be packaged as a jar. Although this directory is a common standard, it will **only** work with war packaging and it will be silently ignored by most build tools if you generate a jar.

Spring Boot also supports advanced resource handling features provided by Spring MVC, allowing use cases such as cache busting static resources or using version agnostic URLs for Webjars.

To use version agnostic URLs for Webjars, simply add the `webjars-locator` dependency. Then declare your Webjar, taking jQuery for example, as `"/webjars/jquery/dist/jquery.min.js"` which results in `"/webjars/jquery/x.y.z/dist/jquery.min.js"` where `x.y.z` is the Webjar version.

Note

If you are using JBoss, you'll need to declare the `webjars-locator-jboss-vfs` dependency instead of the `webjars-locator`; otherwise all Webjars resolve as a 404.

To use cache busting, the following configuration will configure a cache busting solution for all static resources, effectively adding a content hash in URLs, such as `<link href="/css/spring-2a2d595e6ed9a0b24f027f2b63b134d6.css" />`:

```
spring.resources.chain.strategy.content.enabled=true
spring.resources.chain.strategy.content.paths=/**
```

Note

Links to resources are rewritten at runtime in template, thanks to a `ResourceUrlEncodingFilter`, auto-configured for Thymeleaf and FreeMarker. You should manually declare this filter when using JSPs. Other template engines aren't automatically supported right now, but can be with custom template macros/helpers and the use of the [ResourceUrlProvider](#).

When loading resources dynamically with, for example, a JavaScript module loader, renaming files is not an option. That's why other strategies are also supported and can be combined. A "fixed" strategy will add a static version string in the URL, without changing the file name:

```
spring.resources.chain.strategy.content.enabled=true
spring.resources.chain.strategy.content.paths=/**
spring.resources.chain.strategy.fixed.enabled=true
spring.resources.chain.strategy.fixed.paths=/js/lib/
spring.resources.chain.strategy.fixed.version=v12
```

With this configuration, JavaScript modules located under `"/js/lib/"` will use a fixed versioning strategy `"/v12/js/lib/mymodule.js"` while other resources will still use the content one `<link href="/css/spring-2a2d595e6ed9a0b24f027f2b63b134d6.css" />`.

See [ResourceProperties](#) for more of the supported options.

Tip

This feature has been thoroughly described in a dedicated [blog post](#) and in Spring Framework's [reference documentation](#).

Custom Favicon

Spring Boot looks for a `favicon.ico` in the configured static content locations and the root of the classpath (in that order). If such file is present, it is automatically used as the favicon of the application.

ConfigurableWebBindingInitializer

Spring MVC uses a `WebBindingInitializer` to initialize a `WebDataBinder` for a particular request. If you create your own `ConfigurableWebBindingInitializer` @Bean, Spring Boot will automatically configure Spring MVC to use it.

Template engines

As well as REST web services, you can also use Spring MVC to serve dynamic HTML content. Spring MVC supports a variety of templating technologies including Thymeleaf, FreeMarker and JSPs. Many other templating engines also ship their own Spring MVC integrations.

Spring Boot includes auto-configuration support for the following templating engines:

- [FreeMarker](#)
- [Groovy](#)
- [Thymeleaf](#)
- [Mustache](#)

Tip

JSPs should be avoided if possible, there are several [known limitations](#) when using them with embedded servlet containers.

When you're using one of these templating engines with the default configuration, your templates will be picked up automatically from `src/main/resources/templates`.

Tip

IntelliJ IDEA orders the classpath differently depending on how you run your application. Running your application in the IDE via its main method will result in a different ordering to when you run your application using Maven or Gradle or from its packaged jar. This can cause Spring Boot to fail to find the templates on the classpath. If you're affected by this problem you can reorder the classpath in the IDE to place the module's classes and resources first. Alternatively, you can configure the template prefix to search every templates directory on the classpath: `classpath*/templates/`.

Error Handling

Spring Boot provides an `/error` mapping by default that handles all errors in a sensible way, and it is registered as a 'global' error page in the servlet container. For machine clients it will produce a JSON response with details of the error, the HTTP status and the exception message. For browser clients there is a 'whitelabel' error view that renders the same data in HTML format (to customize it just add a `View` that resolves to 'error'). To replace the default behaviour completely you can implement `ErrorController` and register a bean definition of that type, or simply add a bean of type `ErrorAttributes` to use the existing mechanism but replace the contents.

Tip

The `BasicExceptionHandler` can be used as a base class for a custom `ExceptionHandler`. This is particularly useful if you want to add a handler for a new content type (the default is to handle `text/html` specifically and provide a fallback for everything else). To do that just extend `BasicExceptionHandler` and add a public method with a `@RequestMapping` that has a `produces` attribute, and create a bean of your new type.

You can also define a `@ControllerAdvice` to customize the JSON document to return for a particular controller and/or exception type.

```
@ControllerAdvice(basePackageClasses = FooController.class)
public class FooControllerAdvice extends ResponseEntityExceptionHandler {

    @ExceptionHandler(YourException.class)
    @ResponseBody
    ResponseEntity<?> handleControllerException(HttpServletRequest request, Throwable ex) {
        HttpStatus status = getStatus(request);
        return new ResponseEntity<>(new CustomErrorType(status.value(), ex.getMessage()), status);
    }

    private HttpStatus getStatus(HttpServletRequest request) {
        Integer statusCode = (Integer) request.getAttribute("javax.servlet.error.status_code");
        if (statusCode == null) {
            return HttpStatus.INTERNAL_SERVER_ERROR;
        }
        return HttpStatus.valueOf(statusCode);
    }
}
```

In the example above, if `YourException` is thrown by a controller defined in the same package as `FooController`, a json representation of the `CustomErrorType` POJO will be used instead of the `ErrorAttributes` representation.

Custom error pages

If you want to display a custom HTML error page for a given status code, you add a file to an `/error` folder. Error pages can either be static HTML (i.e. added under any of the static resource folders) or built using templates. The name of the file should be the exact status code or a series mask.

For example, to map 404 to a static HTML file, your folder structure would look like this:

```
src/
+- main/
  +- java/
  |   + <source code>
  +- resources/
    +- public/
      +- error/
      |   +- 404.html
      +- <other public assets>
```

To map all 5xx errors using a FreeMarker template, you'd have a structure like this:

```
src/
+- main/
  +- java/
  |   + <source code>
  +- resources/
    +- templates/
      +- error/
      |   +- 5xx.ftl
```

```
+- <other templates>
```

For more complex mappings you can also add beans that implement the `ErrorViewResolver` interface.

```
public class MyErrorViewResolver implements ErrorViewResolver {

    @Override
    public ModelAndView resolveErrorView(HttpServletRequest request,
        HttpStatus status, Map<String, Object> model) {
        // Use the request or status to optionally return a ModelAndView
        return ...
    }

}
```

You can also use regular Spring MVC features like [@ExceptionHandler methods](#) and [@ControllerAdvice](#). The `ErrorController` will then pick up any unhandled exceptions.

Mapping error pages outside of Spring MVC

For applications that aren't using Spring MVC, you can use the `ErrorPageRegistrar` interface to directly register `ErrorPages`. This abstraction works directly with the underlying embedded servlet container and will work even if you don't have a Spring MVC `DispatcherServlet`.

```
@Bean
public ErrorPageRegistrar errorPageRegistrar(){
    return new MyErrorPageRegistrar();
}

// ...

private static class MyErrorPageRegistrar implements ErrorPageRegistrar {

    @Override
    public void registerErrorPages(ErrorPageRegistry registry) {
        registry.addErrorPages(new ErrorPage(HttpStatus.BAD_REQUEST, "/400"));
    }

}
```

N.B. if you register an `ErrorPage` with a path that will end up being handled by a `Filter` (e.g. as is common with some non-Spring web frameworks, like Jersey and Wicket), then the `Filter` has to be explicitly registered as an `ERROR` dispatcher, e.g.

```
@Bean
public FilterRegistrationBean myFilter() {
    FilterRegistrationBean registration = new FilterRegistrationBean();
    registration.setFilter(new MyFilter());
    ...
    registration.setDispatcherTypes(EnumSet.allOf(DispatcherType.class));
    return registration;
}
```

(the default `FilterRegistrationBean` does not include the `ERROR` dispatcher type).

Error Handling on WebSphere Application Server

When deployed to a servlet container, a Spring Boot uses its error page filter to forward a request with an error status to the appropriate error page. The request can only be forwarded to the correct error page if the response has not already been committed. By default, WebSphere Application Server 8.0 and later commits the response upon successful completion of a servlet's service method. You should disable this behaviour by setting `com.ibm.ws.webcontainer.invokeFlushAfterService` to `false`

Spring HATEOAS

If you're developing a RESTful API that makes use of hypermedia, Spring Boot provides auto-configuration for Spring HATEOAS that works well with most applications. The auto-configuration replaces the need to use `@EnableHypermediaSupport` and registers a number of beans to ease building hypermedia-based applications including a `LinkDiscoverers` (for client side support) and an `ObjectMapper` configured to correctly marshal responses into the desired representation. The `ObjectMapper` will be customized based on the `spring.jackson.*` properties or a `Jackson2ObjectMapperBuilder` bean if one exists.

You can take control of Spring HATEOAS's configuration by using `@EnableHypermediaSupport`. Note that this will disable the `ObjectMapper` customization described above.

CORS support

[Cross-origin resource sharing](#) (CORS) is a [W3C specification](#) implemented by [most browsers](#) that allows you to specify in a flexible way what kind of cross domain requests are authorized, instead of using some less secure and less powerful approaches like IFRAME or JSONP.

As of version 4.2, Spring MVC [supports CORS](#) out of the box. Using [controller method CORS configuration](#) with `@CrossOrigin` annotations in your Spring Boot application does not require any specific configuration. [Global CORS configuration](#) can be defined by registering a `WebMvcConfigurer` bean with a customized `addCorsMappings(CorsRegistry)` method:

```
@Configuration
public class MyConfiguration {

    @Bean
    public WebMvcConfigurer corsConfigurer() {
        return new WebMvcConfigurer() {
            @Override
            public void addCorsMappings(CorsRegistry registry) {
                registry.addMapping("/api/**");
            }
        };
    }
}
```

27.2 The ‘Spring WebFlux framework’

Spring WebFlux is the new reactive web framework introduced in Spring Framework 5.0. Unlike Spring MVC, it does not require the Servlet API, is fully asynchronous and non-blocking, and implements the [Reactive Streams](#) specification through [the Reactor project](#).

Spring WebFlux comes in two flavors — the annotation-based one is quite close to the Spring MVC model we know:

```
@RestController
@RequestMapping("/users")
public class MyRestController {

    @GetMapping("/{user}")
    public Mono<User> getUser(@PathVariable Long user) {
        // ...
    }

    @GetMapping("/{user}/customers")
    Flux<Customer> getUserCustomers(@PathVariable Long user) {
        // ...
    }
}
```

```

    }

    @DeleteMapping("/{user}")
    public Mono<User> deleteUser(@PathVariable Long user) {
        // ...
    }
}

```

'WebFlux.fn', the functional variant, separates the routing configuration from the actual handling of the requests:

```

@Configuration
public class RoutingConfiguration {

    @Bean
    public RouterFunction<ServerResponse> monoRouterFunction(UserHandler userHandler) {
        return route(GET("/{user}").and(accept(APPLICATION_JSON)), userHandler::getUser)
            .andRoute(GET("/{user}/customers").and(accept(APPLICATION_JSON)),
                userHandler::getUserCustomers)
            .andRoute(DELETE("/{user}").and(accept(APPLICATION_JSON)), userHandler::deleteUser);
    }
}

@Component
public class UserHandler {

    public Mono<ServerResponse> getUser(ServerRequest request) {
        // ...
    }

    public Mono<ServerResponse> getUserCustomers(ServerRequest request) {
        // ...
    }

    public Mono<ServerResponse> deleteUser(ServerRequest request) {
        // ...
    }
}

```

WebFlux is part of the Spring Framework and detailed information is available in the [reference documentation](#).

To get started, add the `spring-boot-starter-webflux` module to your application.

Note

Adding both `spring-boot-starter-web` and `spring-boot-starter-webflux` modules in your application will result in Spring Boot auto-configuring Spring MVC, not WebFlux. This behavior has been chosen because many Spring developers will add `spring-boot-starter-webflux` to their Spring MVC application to use the reactive `WebClient`. You can still enforce your choice by setting the chosen application type like `SpringApplication.setWebApplicationType(WebApplicationType.REACTIVE)`.

Spring WebFlux auto-configuration

Spring Boot provides auto-configuration for Spring WebFlux that works well with most applications.

The auto-configuration adds the following features on top of Spring's defaults:

- Configuring codecs for `HttpMessageReader` and `HttpMessageWriter` instances (see below).

- Support for serving static resources, including support for WebJars (see below).

If you want to keep Spring Boot WebFlux features, and you just want to add additional [WebFlux configuration](#) you can add your own `@Configuration` class of type `WebFluxConfigurer`, but **without** `@EnableWebFlux`.

If you want to take complete control of Spring WebFlux, you can add your own `@Configuration` annotated with `@EnableWebFlux`.

HTTP codecs with `HttpMessageReaders` and `HttpMessageWriters`

Spring WebFlux uses the `HttpMessageReader` and `HttpMessageWriter` interface to convert HTTP requests and responses. They are configured with `CodecConfigurer` with sensible defaults, by looking at the libraries available in your classpath.

Spring Boot will apply further customization using `CodecCustomizer` instances. For example, `spring.jackson.*` configuration keys will be applied to the Jackson codec.

If you need to add or customize codecs you can create a custom `CodecCustomizer` component:

```
import org.springframework.boot.web.codec.CodecCustomizer;

@Configuration
public class MyConfiguration {

    @Bean
    public CodecCustomizer myCodecCustomizer() {
        return codecConfigurer -> {
            // ...
        }
    }
}
```

You can also leverage [Boot's custom JSON serializers and deserializers](#).

Static Content

By default Spring Boot will serve static content from a directory called `/static` (or `/public` or `/resources` or `/META-INF/resources`) in the classpath. It uses the `ResourceWebHandler` from Spring WebFlux so you can modify that behavior by adding your own `WebFluxConfigurer` and overriding the `addResourceHandlers` method.

By default, resources are mapped on `/**` but you can tune that via `spring.mvc.static-path-pattern`. For instance, relocating all resources to `/resources/**` can be achieved as follows:

```
spring.mvc.static-path-pattern=/resources/**
```

You can also customize the static resource locations using `spring.resources.static-locations` (replacing the default values with a list of directory locations). If you do this the default welcome page detection will switch to your custom locations, so if there is an `index.html` in any of your locations on startup, it will be the home page of the application.

In addition to the 'standard' static resource locations above, a special case is made for [Webjars content](#). Any resources with a path in `/webjars/**` will be served from jar files if they are packaged in the Webjars format.

Tip

Spring WebFlux applications don't strictly depend on the Servlet API, so they can't be deployed as war and have no use of the `src/main/webapp` directory.

Template engines

As well as REST web services, you can also use Spring WebFlux to serve dynamic HTML content. Spring WebFlux supports a variety of templating technologies including Thymeleaf, FreeMarker and Mustache.

Spring Boot includes auto-configuration support for the following templating engines:

- [FreeMarker](#)
- [Thymeleaf](#)
- [Mustache](#)

When you're using one of these templating engines with the default configuration, your templates will be picked up automatically from `src/main/resources/templates`.

27.3 JAX-RS and Jersey

If you prefer the JAX-RS programming model for REST endpoints you can use one of the available implementations instead of Spring MVC. Jersey 1.x and Apache CXF work quite well out of the box if you just register their `Servlet` or `Filter` as a `@Bean` in your application context. Jersey 2.x has some native Spring support so we also provide auto-configuration support for it in Spring Boot together with a starter.

To get started with Jersey 2.x just include the `spring-boot-starter-jersey` as a dependency and then you need one `@Bean` of type `ResourceConfig` in which you register all the endpoints:

```
@Component
public class JerseyConfig extends ResourceConfig {

    public JerseyConfig() {
        register(Endpoint.class);
    }

}
```

Warning

Jersey's support for scanning executable archives is rather limited. For example, it cannot scan for endpoints in a package found in `WEB-INF/classes` when running an executable war file. To avoid this limitation, the `packages` method should not be used and endpoints should be registered individually using the `register` method as shown above.

You can also register an arbitrary number of beans implementing `ResourceConfigCustomizer` for more advanced customizations.

All the registered endpoints should be `@Components` with HTTP resource annotations (`@GET` etc.), e.g.

```
@Component
@Path("/hello")
```

```
public class Endpoint {  
  
    @GET  
    public String message() {  
        return "Hello";  
    }  
  
}
```

Since the `Endpoint` is a Spring `@Component` its lifecycle is managed by Spring and you can `@Autowired` dependencies and inject external configuration with `@Value`. The Jersey servlet will be registered and mapped to `/*` by default. You can change the mapping by adding `@ApplicationPath` to your `ResourceConfig`.

By default Jersey will be set up as a Servlet in a `@Bean` of type `ServletRegistrationBean` named `jerseyServletRegistration`. By default, the servlet will be initialized lazily but you can customize it with `spring.jersey.servlet.load-on-startup`. You can disable or override that bean by creating one of your own with the same name. You can also use a Filter instead of a Servlet by setting `spring.jersey.type=filter` (in which case the `@Bean` to replace or override is `jerseyFilterRegistration`). The servlet has an `@Order` which you can set with `spring.jersey.filter.order`. Both the Servlet and the Filter registrations can be given init parameters using `spring.jersey.init.*` to specify a map of properties.

There is a [Jersey sample](#) so you can see how to set things up. There is also a [Jersey 1.x sample](#). Note that in the Jersey 1.x sample that the spring-boot maven plugin has been configured to unpack some Jersey jars so they can be scanned by the JAX-RS implementation (because the sample asks for them to be scanned in its `Filter` registration). You may need to do the same if any of your JAX-RS resources are packaged as nested jars.

27.4 Embedded servlet container support

Spring Boot includes support for embedded Tomcat, Jetty, and Undertow servers. Most developers will simply use the appropriate 'Starter' to obtain a fully configured instance. By default the embedded server will listen for HTTP requests on port 8080.

Warning

If you choose to use Tomcat on CentOS be aware that, by default, a temporary directory is used to store compiled JSPs, file uploads etc. This directory may be deleted by `tmpwatch` while your application is running leading to failures. To avoid this, you may want to customize your `tmpwatch` configuration so that `tomcat.*` directories are not deleted, or configure `server.tomcat.basedir` so that embedded Tomcat uses a different location.

Servlets, Filters, and listeners

When using an embedded servlet container you can register Servlets, Filters and all the listeners from the Servlet spec (e.g. `HttpSessionListener`) either by using Spring beans or by scanning for Servlet components.

Registering Servlets, Filters, and listeners as Spring beans

Any Servlet, Filter or Servlet `*Listener` instance that is a Spring bean will be registered with the embedded container. This can be particularly convenient if you want to refer to a value from your `application.properties` during configuration.

By default, if the context contains only a single Servlet it will be mapped to `/`. In the case of multiple Servlet beans the bean name will be used as a path prefix. Filters will map to `/*`.

If convention-based mapping is not flexible enough you can use the `ServletRegistrationBean`, `FilterRegistrationBean` and `ServletListenerRegistrationBean` classes for complete control.

Servlet Context Initialization

Embedded servlet containers will not directly execute the Servlet 3.0+ `javax.servlet.ServletContainerInitializer` interface, or Spring's `org.springframework.web.WebApplicationInitializer` interface. This is an intentional design decision intended to reduce the risk that 3rd party libraries designed to run inside a war will break Spring Boot applications.

If you need to perform servlet context initialization in a Spring Boot application, you should register a bean that implements the `org.springframework.boot.web.servlet.ServletContextInitializer` interface. The single `onStartup` method provides access to the `ServletContext`, and can easily be used as an adapter to an existing `WebApplicationInitializer` if necessary.

Scanning for Servlets, Filters, and listeners

When using an embedded container, automatic registration of `@WebServlet`, `@WebFilter`, and `@WebListener` annotated classes can be enabled using `@ServletComponentScan`.

Tip

`@ServletComponentScan` will have no effect in a standalone container, where the container's built-in discovery mechanisms will be used instead.

The ServletWebServerApplicationContext

Under the hood Spring Boot uses a new type of `ApplicationContext` for embedded servlet container support. The `ServletWebServerApplicationContext` is a special type of `WebApplicationContext` that bootstraps itself by searching for a single `ServletWebServerFactory` bean. Usually a `TomcatServletWebServerFactory`, `JettyServletWebServerFactory`, or `UndertowServletWebServerFactory` will have been auto-configured.

Note

You usually won't need to be aware of these implementation classes. Most applications will be auto-configured and the appropriate `ApplicationContext` and `ServletWebServerFactory` will be created on your behalf.

Customizing embedded servlet containers

Common servlet container settings can be configured using Spring `Environment` properties. Usually you would define the properties in your `application.properties` file.

Common server settings include:

- Network settings: listen port for incoming HTTP requests (`server.port`), interface address to bind to `server.address`, etc.
- Session settings: whether the session is persistent (`server.session.persistence`), session timeout (`server.session.timeout`), location of session data (`server.session.store-dir`) and session-cookie configuration (`server.session.cookie.*`).
- Error management: location of the error page (`server.error.path`), etc.
- [SSL](#)
- [HTTP compression](#)

Spring Boot tries as much as possible to expose common settings but this is not always possible. For those cases, dedicated namespaces offer server-specific customizations (see `server.tomcat` and `server.undertow`). For instance, [access logs](#) can be configured with specific features of the embedded servlet container.

Tip

See the [ServerProperties](#) class for a complete list.

Programmatic customization

If you need to configure your embedded servlet container programmatically you can register a Spring bean that implements the `WebServerFactoryCustomizer` interface. `WebServerFactoryCustomizer` provides access to the `ConfigurableServletWebServerFactory` which includes numerous customization setter methods. Dedicated variants exist for Tomcat, Jetty and Undertow.

```
import org.springframework.boot.web.server.WebServerFactoryCustomizer;
import org.springframework.boot.web.servlet.server.ConfigurableServletWebServerFactory;
import org.springframework.stereotype.Component;

@Component
public class CustomizationBean implements
    WebServerFactoryCustomizer<ConfigurableServletWebServerFactory> {

    @Override
    public void customize(ConfigurableServletWebServerFactory server) {
        server.setPort(9000);
    }
}
```

Customizing ConfigurableServletWebServerFactory directly

If the above customization techniques are too limited, you can register the `TomcatServletWebServerFactory`, `JettyServletWebServerFactory` or `UndertowServletWebServerFactory` bean yourself.

```
@Bean
public ConfigurableServletWebServerFactory webServerFactory() {
    TomcatServletWebServerFactory factory = new TomcatServletWebServerFactory();
    factory.setPort(9000);
    factory.setSessionTimeout(10, TimeUnit.MINUTES);
    factory.addErrorPages(new ErrorPage(HttpStatus.NOT_FOUND, "/notfound.html"));
    return factory;
}
```

Setters are provided for many configuration options. Several protected method 'hooks' are also provided should you need to do something more exotic. See the source code documentation for details.

JSP limitations

When running a Spring Boot application that uses an embedded servlet container (and is packaged as an executable archive), there are some limitations in the JSP support.

- With Tomcat it should work if you use war packaging, i.e. an executable war will work, and will also be deployable to a standard container (not limited to, but including Tomcat). An executable jar will not work because of a hard coded file pattern in Tomcat.
- With Jetty it should work if you use war packaging, i.e. an executable war will work, and will also be deployable to any standard container.
- Undertow does not support JSPs.
- Creating a custom `error.jsp` page won't override the default view for [error handling](#), [custom error pages](#) should be used instead.

There is a [JSP sample](#) so you can see how to set things up.

28. Security

If Spring Security is on the classpath then web applications will be secure by default with 'basic' authentication on all HTTP endpoints. To add method-level security to a web application you can also add `@EnableGlobalMethodSecurity` with your desired settings. Additional information can be found in the [Spring Security Reference](#).

The default `AuthenticationManager` has a single user ('user' username and random password, printed at INFO level when the application starts up)

```
Using default security password: 78fa095d-3f4c-48b1-ad50-e24c31d5cf35
```

Note

If you fine-tune your logging configuration, ensure that the `org.springframework.boot.autoconfigure.security` category is set to log INFO messages, otherwise the default password will not be printed.

You can change the password by providing a `security.user.password`. This and other useful properties are externalized via [SecurityProperties](#) (properties prefix "security").

The default security configuration is implemented in `SecurityAutoConfiguration` and in the classes imported from there (`SpringBootWebSecurityConfiguration` for web security and `AuthenticationManagerConfiguration` for authentication configuration which is also relevant in non-web applications). To switch off the default web application security configuration completely you can add a bean with `@EnableWebSecurity` (this does not disable the authentication manager configuration or Actuator's security). To customize it you normally use external properties and beans of type `WebSecurityConfigurerAdapter` (e.g. to add form-based login).

Note

If you add `@EnableWebSecurity` and also disable Actuator security, you will get the default form-based login for the entire application unless you add a custom `WebSecurityConfigurerAdapter`.

To also switch off the authentication manager configuration you can add a bean of type `AuthenticationManager`, or else configure the global `AuthenticationManager` by autowiring an `AuthenticationManagerBuilder` into a method in one of your `@Configuration` classes. There are several secure applications in the [Spring Boot samples](#) to get you started with common use cases.

The basic features you get out of the box in a web application are:

- An `AuthenticationManager` bean with in-memory store and a single user (see `SecurityProperties.User` for the properties of the user).
- Ignored (insecure) paths for common static resource locations (`/css/**`, `/js/**`, `/images/**`, `/webjars/**` and `**/favicon.ico`).
- HTTP Basic security for all other endpoints.
- Security events published to Spring's `ApplicationEventPublisher` (successful and unsuccessful authentication and access denied).

- Common low-level features (HSTS, XSS, CSRF, caching) provided by Spring Security are on by default.

All of the above can be switched on and off or modified using external properties (`security.*`). To override the access rules without changing any other auto-configured features add a `@Bean` of type `WebSecurityConfigurerAdapter` with `@Order(SecurityProperties.ACCESS_OVERRIDE_ORDER)` and configure it to meet your needs.

Note

By default, a `WebSecurityConfigurerAdapter` will match any path. If you don't want to completely override Spring Boot's auto-configured access rules, your adapter must explicitly configure the paths that you do want to override.

28.1 OAuth2

If you have `spring-security-oauth2` on your classpath you can take advantage of some auto-configuration to make it easy to set up Authorization or Resource Server. For full details, see the [Spring Security OAuth 2 Developers Guide](#).

Authorization Server

To create an Authorization Server and grant access tokens you need to use `@EnableAuthorizationServer` and provide `security.oauth2.client.client-id` and `security.oauth2.client.client-secret` properties. The client will be registered for you in an in-memory repository.

Having done that you will be able to use the client credentials to create an access token, for example:

```
$ curl client:secret@localhost:8080/oauth/token -d grant_type=password -d username=user -d password=pwd
```

The basic auth credentials for the `/token` endpoint are the `client-id` and `client-secret`. The user credentials are the normal Spring Security user details (which default in Spring Boot to “user” and a random password).

To switch off the auto-configuration and configure the Authorization Server features yourself just add a `@Bean` of type `AuthorizationServerConfigurer`.

Resource Server

To use the access token you need a Resource Server (which can be the same as the Authorization Server). Creating a Resource Server is easy, just add `@EnableResourceServer` and provide some configuration to allow the server to decode access tokens. If your application is also an Authorization Server it already knows how to decode tokens, so there is nothing else to do. If your app is a standalone service then you need to give it some more configuration, one of the following options:

- `security.oauth2.resource.user-info-uri` to use the `/me` resource (e.g. `https://uaa.run.pivotal.io/userinfo` on PWS)
- `security.oauth2.resource.token-info-uri` to use the token decoding endpoint (e.g. `https://uaa.run.pivotal.io/check_token` on PWS).

If you specify both the `user-info-uri` and the `token-info-uri` then you can set a flag to say that one is preferred over the other (`prefer-token-info=true` is the default).

Alternatively (instead of `user-info-uri` or `token-info-uri`) if the tokens are JWTs you can configure a `security.oauth2.resource.jwt.key-value` to decode them locally (where the key is a verification key). The verification key value is either a symmetric secret or PEM-encoded RSA public key. If you don't have the key and it's public you can provide a URI where it can be downloaded (as a JSON object with a "value" field) with `security.oauth2.resource.jwt.key-uri`. E.g. on PWS:

```
$ curl https://uaa.run.pivotal.io/token_key
{"alg":"SHA256withRSA","value":"-----BEGIN PUBLIC KEY-----\nMIIBI...\n-----END PUBLIC KEY-----\n"}
```

Warning

If you use the `security.oauth2.resource.jwt.key-uri` the authorization server needs to be running when your application starts up. It will log a warning if it can't find the key, and tell you what to do to fix it.

OAuth2 resources are protected by a filter chain with order `security.oauth2.resource.filter-order` and the default is after the filter protecting the actuator endpoints by default (so actuator endpoints will stay on HTTP Basic unless you change the order).

28.2 Token Type in User Info

Google, and certain other 3rd party identity providers, are more strict about the token type name that is sent in the headers to the user info endpoint. The default is "Bearer" which suits most providers and matches the spec, but if you need to change it you can set `security.oauth2.resource.token-type`.

28.3 Customizing the User Info RestTemplate

If you have a `user-info-uri`, the resource server features use an `OAuth2RestTemplate` internally to fetch user details for authentication. This is provided as a `@Bean` of type `UserInfoRestTemplateFactory`. The default should be fine for most providers, but occasionally you might need to add additional interceptors, or change the request authenticator (which is how the token gets attached to outgoing requests). To add a customization just create a bean of type `UserInfoRestTemplateCustomizer` - it has a single method that will be called after the bean is created but before it is initialized. The rest template that is being customized here is *only* used internally to carry out authentication. Alternatively, you could define your own `UserInfoRestTemplateFactory` `@Bean` to take full control.

Tip

To set an RSA key value in YAML use the "pipe" continuation marker to split it over multiple lines ("|") and remember to indent the key value (it's a standard YAML language feature). Example:

```
security:
  oauth2:
    resource:
      jwt:
        keyValue: |
          -----BEGIN PUBLIC KEY-----
          MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKC...
          -----END PUBLIC KEY-----
```

Client

To make your web-app into an OAuth2 client you can simply add `@EnableOAuth2Client` and Spring Boot will create a `OAuth2ClientContext` and `OAuth2ProtectedResourceDetails` that are necessary to create an `OAuth2RestOperations`. Spring Boot does not automatically create such bean but you can easily create your own:

```
@Bean
public OAuth2RestTemplate oauth2RestTemplate(OAuth2ClientContext oauth2ClientContext,
    OAuth2ProtectedResourceDetails details) {
    return new OAuth2RestTemplate(details, oauth2ClientContext);
}
```

Note

You may want to add a qualifier and review your configuration as more than one `RestTemplate` may be defined in your application.

This configuration uses `security.oauth2.client.*` as credentials (the same as you might be using in the Authorization Server), but in addition it will need to know the authorization and token URIs in the Authorization Server. For example:

application.yml.

```
security:
  oauth2:
    client:
      clientId: bd1c0a783ccdd1c9b9e4
      clientSecret: 1a9030fbca47a5b2c28e92f19050bb77824b5ad1
      accessTokenUri: https://github.com/login/oauth/access_token
      userAuthorizationUri: https://github.com/login/oauth/authorize
      clientAuthenticationScheme: form
```

An application with this configuration will redirect to Github for authorization when you attempt to use the `OAuth2RestTemplate`. If you are already signed into Github you won't even notice that it has authenticated. These specific credentials will only work if your application is running on port 8080 (register your own client app in Github or other provider for more flexibility).

To limit the scope that the client asks for when it obtains an access token you can set `security.oauth2.client.scope` (comma separated or an array in YAML). By default the scope is empty and it is up to Authorization Server to decide what the defaults should be, usually depending on the settings in the client registration that it holds.

Note

There is also a setting for `security.oauth2.client.client-authentication-scheme` which defaults to "header" (but you might need to set it to "form" if, like Github for instance, your OAuth2 provider doesn't like header authentication). In fact, the `security.oauth2.client.*` properties are bound to an instance of `AuthorizationCodeResourceDetails` so all its properties can be specified.

Tip

In a non-web application you can still create an `OAuth2RestOperations` and it is still wired into the `security.oauth2.client.*` configuration. In this case it is a "client credentials token

grant” you will be asking for if you use it (and there is no need to use `@EnableOAuth2Client` or `@EnableOAuth2Sso`). To prevent that infrastructure to be defined, just remove the `security.oauth2.client.client-id` from your configuration (or make it the empty string).

Single Sign On

An OAuth2 Client can be used to fetch user details from the provider (if such features are available) and then convert them into an `Authentication` token for Spring Security. The Resource Server above support this via the `user-info-uri` property This is the basis for a Single Sign On (SSO) protocol based on OAuth2, and Spring Boot makes it easy to participate by providing an annotation `@EnableOAuth2Sso`. The Github client above can protect all its resources and authenticate using the Github `/user/` endpoint, by adding that annotation and declaring where to find the endpoint (in addition to the `security.oauth2.client.*` configuration already listed above):

application.yml.

```
security:
  oauth2:
  ...
  resource:
    userInfoUri: https://api.github.com/user
    preferTokenInfo: false
```

Since all paths are secure by default, there is no “home” page that you can show to unauthenticated users and invite them to login (by visiting the `/login` path, or the path specified by `security.oauth2.sso.login-path`).

To customize the access rules or paths to protect, so you can add a “home” page for instance, `@EnableOAuth2Sso` can be added to a `WebSecurityConfigurerAdapter` and the annotation will cause it to be decorated and enhanced with the necessary pieces to get the `/login` path working. For example, here we simply allow unauthenticated access to the home page at `/` and keep the default for everything else:

```
@Configuration
static class WebSecurityConfiguration extends WebSecurityConfigurerAdapter {

    @Override
    public void init(WebSecurity web) {
        web.ignoring().antMatchers("/");
    }

    @Override
    protected void configure(HttpSecurity http) throws Exception {
        http.antMatcher("/**").authorizeRequests().anyRequest().authenticated();
    }

}
```

28.4 Actuator Security

If the Actuator is also in use, you will find:

- The management endpoints are secure even if the application endpoints are insecure.
- Security events are transformed into `AuditEvent` instances and published to the `AuditEventRepository`.
- The default user will have the `ACTUATOR` role as well as the `USER` role.

The Actuator security features can be modified using external properties (`management.security.*`). To override the application access rules add a `@Bean` of type `WebSecurityConfigurerAdapter` and use `@Order(SecurityProperties.ACCESS_OVERRIDE_ORDER)` if you *don't* want to override the actuator access rules, or `@Order(ManagementServerProperties.ACCESS_OVERRIDE_ORDER)` if you *do* want to override the actuator access rules.

29. Working with SQL databases

The Spring Framework provides extensive support for working with SQL databases. From direct JDBC access using `JdbcTemplate` to complete ‘object relational mapping’ technologies such as Hibernate. Spring Data provides an additional level of functionality, creating `Repository` implementations directly from interfaces and using conventions to generate queries from your method names.

29.1 Configure a DataSource

Java’s `javax.sql.DataSource` interface provides a standard method of working with database connections. Traditionally a `DataSource` uses a URL along with some credentials to establish a database connection.

Tip

Check also [the ‘How-to’ section](#) for more advanced examples, typically to take full control over the configuration of the `DataSource`.

Embedded Database Support

It’s often convenient to develop applications using an in-memory embedded database. Obviously, in-memory databases do not provide persistent storage; you will need to populate your database when your application starts and be prepared to throw away data when your application ends.

Tip

The ‘How-to’ section includes a [section on how to initialize a database](#)

Spring Boot can auto-configure embedded [H2](#), [HSQL](#) and [Derby](#) databases. You don’t need to provide any connection URLs, simply include a build dependency to the embedded database that you want to use.

Note

If you are using this feature in your tests, you may notice that the same database is reused by your whole test suite regardless of the number of application contexts that you use. If you want to make sure that each context has a separate embedded database, you should set `spring.datasource.generate-unique-name` to `true`.

For example, typical POM dependencies would be:

```
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-data-jpa</artifactId>
</dependency>
<dependency>
  <groupId>org.hsqldb</groupId>
  <artifactId>hsqldb</artifactId>
  <scope>runtime</scope>
</dependency>
```

Note

You need a dependency on `spring-jdbc` for an embedded database to be auto-configured. In this example it's pulled in transitively via `spring-boot-starter-data-jpa`.

Tip

If, for whatever reason, you do configure the connection URL for an embedded database, care should be taken to ensure that the database's automatic shutdown is disabled. If you're using H2 you should use `DB_CLOSE_ON_EXIT=FALSE` to do so. If you're using HSQLDB, you should ensure that `shutdown=true` is not used. Disabling the database's automatic shutdown allows Spring Boot to control when the database is closed, thereby ensuring that it happens once access to the database is no longer needed.

Connection to a production database

Production database connections can also be auto-configured using a pooling `DataSource`. Here's the algorithm for choosing a specific implementation:

- We prefer HikariCP for its performance and concurrency, so if that is available we always choose it.
- Otherwise, if the Tomcat pooling `DataSource` is available we will use it.
- If neither HikariCP nor the Tomcat pooling `datasource` are available and if Commons DBCP2 is available we will use it.

If you use the `spring-boot-starter-jdbc` or `spring-boot-starter-data-jpa` 'starters' you will automatically get a dependency to HikariCP.

Note

You can bypass that algorithm completely and specify the connection pool to use via the `spring.datasource.type` property. This is especially important if you are running your application in a Tomcat container as `tomcat-jdbc` is provided by default.

Tip

Additional connection pools can always be configured manually. If you define your own `DataSource` bean, auto-configuration will not occur.

`DataSource` configuration is controlled by external configuration properties in `spring.datasource.*`. For example, you might declare the following section in `application.properties`:

```
spring.datasource.url=jdbc:mysql://localhost/test
spring.datasource.username=dbuser
spring.datasource.password=dbpass
spring.datasource.driver-class-name=com.mysql.jdbc.Driver
```

Note

You should at least specify the url using the `spring.datasource.url` property or Spring Boot will attempt to auto-configure an embedded database.

Tip

You often won't need to specify the `driver-class-name` since Spring boot can deduce it for most databases from the `url`.

Note

For a pooling `DataSource` to be created we need to be able to verify that a valid `Driver` class is available, so we check for that before doing anything. I.e. if you set `spring.datasource.driver-class-name=com.mysql.jdbc.Driver` then that class has to be loadable.

See [DataSourceProperties](#) for more of the supported options. These are the standard options that work regardless of the actual implementation. It is also possible to fine-tune implementation-specific settings using their respective prefix (`spring.datasource.hikari.*`, `spring.datasource.tomcat.*`, and `spring.datasource.dbcp2.*`). Refer to the documentation of the connection pool implementation you are using for more details.

For instance, if you are using the [Tomcat connection pool](#) you could customize many additional settings:

```
# Number of ms to wait before throwing an exception if no connection is available.
spring.datasource.tomcat.max-wait=10000

# Maximum number of active connections that can be allocated from this pool at the same time.
spring.datasource.tomcat.max-active=50

# Validate the connection before borrowing it from the pool.
spring.datasource.tomcat.test-on-borrow=true
```

Connection to a JNDI DataSource

If you are deploying your Spring Boot application to an Application Server you might want to configure and manage your `DataSource` using your Application Servers built-in features and access it using JNDI.

The `spring.datasource.jndi-name` property can be used as an alternative to the `spring.datasource.url`, `spring.datasource.username` and `spring.datasource.password` properties to access the `DataSource` from a specific JNDI location. For example, the following section in `application.properties` shows how you can access a JBoss AS defined `DataSource`:

```
spring.datasource.jndi-name=java:jboss/datasources/customers
```

29.2 Using JdbcTemplate

Spring's `JdbcTemplate` and `NamedParameterJdbcTemplate` classes are auto-configured and you can `@Autowired` them directly into your own beans:

```
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.jdbc.core.JdbcTemplate;
import org.springframework.stereotype.Component;

@Component
public class MyBean {

    private final JdbcTemplate jdbcTemplate;

    @Autowired
```

```

public MyBean(JdbcTemplate jdbcTemplate) {
    this.jdbcTemplate = jdbcTemplate;
}

// ...
}

```

You can customize some properties of the template using the `spring.jdbc.template.*` properties:

```
spring.jdbc.template.max-rows=500
```

Tip

The `NamedParameterJdbcTemplate` reuses the same `JdbcTemplate` instance behind the scene. If more than one `JdbcTemplate` is defined and no primary candidate exists, the `NamedParameterJdbcTemplate` is not auto-configured.

29.3 JPA and ‘Spring Data’

The Java Persistence API is a standard technology that allows you to ‘map’ objects to relational databases. The `spring-boot-starter-data-jpa` POM provides a quick way to get started. It provides the following key dependencies:

- Hibernate — One of the most popular JPA implementations.
- Spring Data JPA — Makes it easy to implement JPA-based repositories.
- Spring ORMs — Core ORM support from the Spring Framework.

Tip

We won’t go into too many details of JPA or Spring Data here. You can follow the [‘Accessing Data with JPA’](#) guide from [spring.io](#) and read the [Spring Data JPA](#) and [Hibernate](#) reference documentation.

Entity Classes

Traditionally, JPA ‘Entity’ classes are specified in a `persistence.xml` file. With Spring Boot this file is not necessary and instead ‘Entity Scanning’ is used. By default all packages below your main configuration class (the one annotated with `@EnableAutoConfiguration` or `@SpringBootApplication`) will be searched.

Any classes annotated with `@Entity`, `@Embeddable` or `@MappedSuperclass` will be considered. A typical entity class would look something like this:

```

package com.example.myapp.domain;

import java.io.Serializable;
import javax.persistence.*;

@Entity
public class City implements Serializable {

    @Id
    @GeneratedValue
    private Long id;
}

```

```

@Column(nullable = false)
private String name;

@Column(nullable = false)
private String state;

// ... additional members, often include @OneToMany mappings

protected City() {
    // no-args constructor required by JPA spec
    // this one is protected since it shouldn't be used directly
}

public City(String name, String state) {
    this.name = name;
    this.country = country;
}

public String getName() {
    return this.name;
}

public String getState() {
    return this.state;
}

// ... etc
}

```

Tip

You can customize entity scanning locations using the `@EntityScan` annotation. See the [Section 78.4, “Separate @Entity definitions from Spring configuration”](#) how-to.

Spring Data JPA Repositories

Spring Data JPA repositories are interfaces that you can define to access data. JPA queries are created automatically from your method names. For example, a `CityRepository` interface might declare a `findAllByState(String state)` method to find all cities in a given state.

For more complex queries you can annotate your method using Spring Data’s [Query](#) annotation.

Spring Data repositories usually extend from the [Repository](#) or [CrudRepository](#) interfaces. If you are using auto-configuration, repositories will be searched from the package containing your main configuration class (the one annotated with `@EnableAutoConfiguration` or `@SpringBootApplication`) down.

Here is a typical Spring Data repository:

```

package com.example.myapp.domain;

import org.springframework.data.domain.*;
import org.springframework.data.repository.*;

public interface CityRepository extends Repository<City, Long> {

    Page<City> findAll(Pageable pageable);

    City findByNameAndCountryAllIgnoringCase(String name, String country);

}

```

Tip

We have barely scratched the surface of Spring Data JPA. For complete details check their [reference documentation](#).

Creating and dropping JPA databases

By default, JPA databases will be automatically created **only** if you use an embedded database (H2, HSQL or Derby). You can explicitly configure JPA settings using `spring.jpa.*` properties. For example, to create and drop tables you can add the following to your `application.properties`.

```
spring.jpa.hibernate.ddl-auto=create-drop
```

Note

Hibernate's own internal property name for this (if you happen to remember it better) is `hibernate.hbm2ddl.auto`. You can set it, along with other Hibernate native properties, using `spring.jpa.properties.*` (the prefix is stripped before adding them to the entity manager). Example:

```
spring.jpa.properties.hibernate.globally_quoted_identifiers=true
```

passes `hibernate.globally_quoted_identifiers` to the Hibernate entity manager.

By default the DDL execution (or validation) is deferred until the `ApplicationContext` has started. There is also a `spring.jpa.generate-ddl` flag, but it is not used if Hibernate autoconfig is active because the `ddl-auto` settings are more fine-grained.

Open EntityManager in View

If you are running a web application, Spring Boot will by default register [OpenEntityManagerInViewInterceptor](#) to apply the "Open EntityManager in View" pattern, i.e. to allow for lazy loading in web views. If you don't want this behavior you should set `spring.jpa.open-in-view` to `false` in your `application.properties`.

29.4 Using H2's web console

The [H2 database](#) provides a [browser-based console](#) that Spring Boot can auto-configure for you. The console will be auto-configured when the following conditions are met:

- You are developing a web application
- `com.h2database:h2` is on the classpath
- You are using [Spring Boot's developer tools](#)

Tip

If you are not using Spring Boot's developer tools, but would still like to make use of H2's console, then you can do so by configuring the `spring.h2.console.enabled` property with a value of `true`. The H2 console is only intended for use during development so care should be taken to ensure that `spring.h2.console.enabled` is not set to `true` in production.

Changing the H2 console's path

By default the console will be available at `/h2-console`. You can customize the console's path using the `spring.h2.console.path` property.

Securing the H2 console

When Spring Security is on the classpath and basic auth is enabled, the H2 console will be automatically secured using basic auth. The following properties can be used to customize the security configuration:

- `security.user.role`
- `security.basic.authorize-mode`
- `security.basic.enabled`

29.5 Using jOOQ

Java Object Oriented Querying ([jOOQ](#)) is a popular product from [Data Geekery](#) which generates Java code from your database, and lets you build type safe SQL queries through its fluent API. Both the commercial and open source editions can be used with Spring Boot.

Code Generation

In order to use jOOQ type-safe queries, you need to generate Java classes from your database schema. You can follow the instructions in the [jOOQ user manual](#). If you are using the `jooq-codegen-maven` plugin (and you also use the `spring-boot-starter-parent` “parent POM”) you can safely omit the plugin's `<version>` tag. You can also use Spring Boot defined version variables (e.g. `h2.version`) to declare the plugin's database dependency. Here's an example:

```
<plugin>
  <groupId>org.jooq</groupId>
  <artifactId>jooq-codegen-maven</artifactId>
  <executions>
    ...
  </executions>
  <dependencies>
    <dependency>
      <groupId>com.h2database</groupId>
      <artifactId>h2</artifactId>
      <version>${h2.version}</version>
    </dependency>
  </dependencies>
  <configuration>
    <jdbc>
      <driver>org.h2.Driver</driver>
      <url>jdbc:h2:~/yourdatabase</url>
    </jdbc>
    <generator>
      ...
    </generator>
  </configuration>
</plugin>
```

Using DSLContext

The fluent API offered by jOOQ is initiated via the `org.jooq.DSLContext` interface. Spring Boot will auto-configure a `DSLContext` as a Spring Bean and connect it to your application `DataSource`. To use the `DSLContext` you can just `@Autowired` it:

```
@Component
public class JooqExample implements CommandLineRunner {

    private final DSLContext create;

    @Autowired
    public JooqExample(DSLContext dslContext) {
        this.create = dslContext;
    }
}
```

Tip

The jOOQ manual tends to use a variable named `create` to hold the `DSLContext`, we've done the same for this example.

You can then use the `DSLContext` to construct your queries:

```
public List<GregorianCalendar> authorsBornAfter1980() {
    return this.create.selectFrom(AUTHOR)
        .where(AUTHOR.DATE_OF_BIRTH.greaterThan(new GregorianCalendar(1980, 0, 1)))
        .fetch(AUTHOR.DATE_OF_BIRTH);
}
```

jOOQ SQL dialect

Spring Boot determines the SQL dialect to use for your datasource unless the `spring.jooq.sql-dialect` property has been configured. If the dialect couldn't be detected, `DEFAULT` is used.

Note

Spring Boot can only auto-configure dialects supported by the open source version of jOOQ.

Customizing jOOQ

More advanced customizations can be achieved by defining your own `@Bean` definitions which will be used when the `jOOQ Configuration` is created. You can define beans for the following jOOQ Types:

- `ConnectionProvider`
- `TransactionProvider`
- `RecordMapperProvider`
- `RecordListenerProvider`
- `ExecuteListenerProvider`
- `VisitListenerProvider`

You can also create your own `org.jooq.Configuration @Bean` if you want to take complete control of the jOOQ configuration.

30. Working with NoSQL technologies

Spring Data provides additional projects that help you access a variety of NoSQL technologies including [MongoDB](#), [Neo4J](#), [Elasticsearch](#), [Solr](#), [Redis](#), [Gemfire](#), [Cassandra](#), [Couchbase](#) and [LDAP](#). Spring Boot provides auto-configuration for Redis, MongoDB, Neo4j, Elasticsearch, Solr, Cassandra, Couchbase and LDAP; you can make use of the other projects, but you will need to configure them yourself. Refer to the appropriate reference documentation at projects.spring.io/spring-data.

30.1 Redis

[Redis](#) is a cache, message broker and richly-featured key-value store. Spring Boot offers basic auto-configuration for the [Jedis](#) and [Lettuce](#) client library and abstractions on top of it provided by [Spring Data Redis](#).

There is a `spring-boot-starter-data-redis` ‘Starter’ for collecting the dependencies in a convenient way that uses [Jedis](#) by default. If you are building a reactive application, the `spring-boot-starter-data-redis-reactive` ‘Starter’ will get you going.

Connecting to Redis

You can inject an auto-configured `RedisConnectionFactory`, `StringRedisTemplate` or vanilla `RedisTemplate` instance as you would any other Spring Bean. By default the instance will attempt to connect to a Redis server using `localhost:6379`:

```
@Component
public class MyBean {

    private StringRedisTemplate template;

    @Autowired
    public MyBean(StringRedisTemplate template) {
        this.template = template;
    }

    // ...
}
```

Tip

You can also register an arbitrary number of beans implementing `JedisClientConfigurationBuilderCustomizer` for more advanced customizations. If you are using [Lettuce](#), `LettuceClientConfigurationBuilderCustomizer` is also available.

If you add a `@Bean` of your own of any of the auto-configured types it will replace the default (except in the case of `RedisTemplate` the exclusion is based on the bean name ‘`redisTemplate`’ not its type). If `commons-pool2` is on the classpath you will get a pooled connection factory by default.

30.2 MongoDB

[MongoDB](#) is an open-source NoSQL document database that uses a JSON-like schema instead of traditional table-based relational data. Spring Boot offers several conveniences for working with MongoDB, including the `spring-boot-starter-data-mongodb` and `spring-boot-starter-data-mongodb-reactive` ‘Starters’.

Connecting to a MongoDB database

You can inject an auto-configured `org.springframework.data.mongodb.MongoDbFactory` to access Mongo databases. By default the instance will attempt to connect to a MongoDB server using the URL `mongodb://localhost/test`:

```
import org.springframework.data.mongodb.MongoDbFactory;
import com.mongodb.DB;

@Component
public class MyBean {

    private final MongoDbFactory mongo;

    @Autowired
    public MyBean(MongoDbFactory mongo) {
        this.mongo = mongo;
    }

    // ...

    public void example() {
        DB db = mongo.getDb();
        // ...
    }
}
```

You can set `spring.data.mongodb.uri` property to change the URL and configure additional settings such as the *replica set*:

```
spring.data.mongodb.uri=mongodb://user:secret@mongo1.example.com:12345,mongo2.example.com:23456/test
```

Alternatively, as long as you're using Mongo 2.x, specify a *host/port*. For example, you might declare the following in your `application.properties`:

```
spring.data.mongodb.host=mongoserver
spring.data.mongodb.port=27017
```

Note

`spring.data.mongodb.host` and `spring.data.mongodb.port` are not supported if you're using the Mongo 3.0 Java driver. In such cases, `spring.data.mongodb.uri` should be used to provide all of the configuration.

Tip

If `spring.data.mongodb.port` is not specified the default of `27017` is used. You could simply delete this line from the sample above.

Tip

If you aren't using Spring Data Mongo you can inject `com.mongodb.Mongo` beans instead of using `MongoDbFactory`.

You can also declare your own `MongoDbFactory` or `Mongo` bean if you want to take complete control of establishing the MongoDB connection.

MongoTemplate

Spring Data Mongo provides a [MongoTemplate](#) class that is very similar in its design to Spring's `JdbcTemplate`. As with `JdbcTemplate` Spring Boot auto-configures a bean for you to simply inject:

```
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.data.mongodb.core.MongoTemplate;
import org.springframework.stereotype.Component;

@Component
public class MyBean {

    private final MongoTemplate mongoTemplate;

    @Autowired
    public MyBean(MongoTemplate mongoTemplate) {
        this.mongoTemplate = mongoTemplate;
    }

    // ...
}
```

See the `MongoOperations` Javadoc for complete details.

Spring Data MongoDB repositories

Spring Data includes repository support for MongoDB. As with the JPA repositories discussed earlier, the basic principle is that queries are constructed for you automatically based on method names.

In fact, both Spring Data JPA and Spring Data MongoDB share the same common infrastructure; so you could take the JPA example from earlier and, assuming that `City` is now a Mongo data class rather than a JPA `@Entity`, it will work in the same way.

```
package com.example.myapp.domain;

import org.springframework.data.domain.*;
import org.springframework.data.repository.*;

public interface CityRepository extends Repository<City, Long> {

    Page<City> findAll(Pageable pageable);

    City findByNameAndCountryAllIgnoringCase(String name, String country);
}
```

Tip

For complete details of Spring Data MongoDB, including its rich object mapping technologies, refer to their [reference documentation](#).

Embedded Mongo

Spring Boot offers auto-configuration for [Embedded Mongo](#). To use it in your Spring Boot application add a dependency on `de.flapdoodle.embed:de.flapdoodle.embed.mongo`.

The port that Mongo will listen on can be configured using the `spring.data.mongodb.port` property. To use a randomly allocated free port use a value of zero. The `MongoClient` created by `MongoAutoConfiguration` will be automatically configured to use the randomly allocated port.

Note

If you do not configure a custom port, the embedded support will use a random port by default (rather than 27017).

If you have SLF4J on the classpath, output produced by Mongo will be automatically routed to a logger named `org.springframework.boot.autoconfigure.mongo.embedded.EmbeddedMongo`.

You can declare your own `IMongodConfig` and `IRuntimeConfig` beans to take control of the Mongo instance's configuration and logging routing.

30.3 Neo4j

[Neo4j](#) is an open-source NoSQL graph database that uses a rich data model of nodes related by first class relationships which is better suited for connected big data than traditional rdbms approaches. Spring Boot offers several conveniences for working with Neo4j, including the `spring-boot-starter-data-neo4j` 'Starter'.

Connecting to a Neo4j database

You can inject an auto-configured `Neo4jSession`, `Session` or `Neo4jOperations` instance as you would any other Spring Bean. By default the instance will attempt to connect to a Neo4j server using `localhost:7474`:

```
@Component
public class MyBean {

    private final Neo4jTemplate neo4jTemplate;

    @Autowired
    public MyBean(Neo4jTemplate neo4jTemplate) {
        this.neo4jTemplate = neo4jTemplate;
    }

    // ...

}
```

You can take full control of the configuration by adding a `org.neo4j.ogm.config.Configuration` `@Bean` of your own. Also, adding a `@Bean` of type `Neo4jOperations` disables the auto-configuration.

You can configure the user and credentials to use via the `spring.data.neo4j.*` properties:

```
spring.data.neo4j.uri=http://my-server:7474
spring.data.neo4j.username=neo4j
spring.data.neo4j.password=secret
```

Using the embedded mode

If you add `org.neo4j:neo4j-ogm-embedded-driver` to the dependencies of your application, Spring Boot will automatically configure an in-process embedded instance of Neo4j that will not persist any data when your application shuts down. You can explicitly disable that mode using `spring.data.neo4j.embedded.enabled=false`. You can also enable persistence for the embedded mode:

```
spring.data.neo4j.uri=file://var/tmp/graph.db
```

Note

The Neo4j OGM embedded driver does not provide the Neo4j kernel. Users are expected to provide this dependency manually, see [the documentation](#) for more details.

Neo4jSession

By default, if you are running a web application, the session is bound to the thread for the entire processing of the request (i.e. the "Open Session in View" pattern). If you don't want this behavior add the following to your `application.properties`:

```
spring.data.neo4j.open-in-view=false
```

Spring Data Neo4j repositories

Spring Data includes repository support for Neo4j.

In fact, both Spring Data JPA and Spring Data Neo4j share the same common infrastructure; so you could take the JPA example from earlier and, assuming that `City` is now a Neo4j OGM `@NodeEntity` rather than a JPA `@Entity`, it will work in the same way.

Tip

You can customize entity scanning locations using the `@EntityScan` annotation.

To enable repository support (and optionally support for `@Transactional`), add the following two annotations to your Spring configuration:

```
@EnableNeo4jRepositories(basePackages = "com.example.myapp.repository")
@EnableTransactionManagement
```

Repository example

```
package com.example.myapp.domain;

import org.springframework.data.domain.*;
import org.springframework.data.repository.*;

public interface CityRepository extends GraphRepository<City> {

    Page<City> findAll(Pageable pageable);

    City findByNameAndCountry(String name, String country);

}
```

Tip

For complete details of Spring Data Neo4j, including its rich object mapping technologies, refer to their [reference documentation](#).

30.4 Gemfire

[Spring Data Gemfire](#) provides convenient Spring-friendly tools for accessing the [Pivotal Gemfire](#) data management platform. There is a `spring-boot-starter-data-gemfire` 'Starter' for collecting the

dependencies in a convenient way. There is currently no auto-configuration support for Gemfire, but you can enable Spring Data Repositories with a [single annotation](#) (`@EnableGemfireRepositories`).

30.5 Solr

[Apache Solr](#) is a search engine. Spring Boot offers basic auto-configuration for the Solr 5 client library and abstractions on top of it provided by [Spring Data Solr](#). There is a `spring-boot-starter-data-solr` ‘Starter’ for collecting the dependencies in a convenient way.

Connecting to Solr

You can inject an auto-configured `SolrClient` instance as you would any other Spring bean. By default the instance will attempt to connect to a server using `localhost:8983/solr`:

```
@Component
public class MyBean {

    private SolrClient solr;

    @Autowired
    public MyBean(SolrClient solr) {
        this.solr = solr;
    }

    // ...
}
```

If you add a `@Bean` of your own of type `SolrClient` it will replace the default.

Spring Data Solr repositories

Spring Data includes repository support for Apache Solr. As with the JPA repositories discussed earlier, the basic principle is that queries are constructed for you automatically based on method names.

In fact, both Spring Data JPA and Spring Data Solr share the same common infrastructure; so you could take the JPA example from earlier and, assuming that `City` is now a `@SolrDocument` class rather than a JPA `@Entity`, it will work in the same way.

Tip

For complete details of Spring Data Solr, refer to their [reference documentation](#).

30.6 Elasticsearch

[Elasticsearch](#) is an open source, distributed, real-time search and analytics engine. Spring Boot offers basic auto-configuration for the Elasticsearch and abstractions on top of it provided by [Spring Data Elasticsearch](#). There is a `spring-boot-starter-data-elasticsearch` ‘Starter’ for collecting the dependencies in a convenient way. Spring Boot also supports [Jest](#).

Connecting to Elasticsearch using Jest

If you have `Jest` on the classpath, you can inject an auto-configured `JestClient` targeting `localhost:9200` by default. You can further tune how the client is configured:

```
spring.elasticsearch.jest.uris=http://search.example.com:9200
spring.elasticsearch.jest.read-timeout=10000
spring.elasticsearch.jest.username=user
```

```
spring.elasticsearch.jest.password=secret
```

You can also register an arbitrary number of beans implementing `HttpClientConfigBuilderCustomizer` for more advanced customizations. The example below tunes additional HTTP settings:

```
static class HttpSettingsCustomizer implements HttpClientConfigBuilderCustomizer {

    @Override
    public void customize(HttpClientConfig.Builder builder) {
        builder.maxTotalConnection(100).defaultMaxTotalConnectionPerRoute(5);
    }
}
```

To take full control over the registration, define a `JestClient` bean.

Connecting to Elasticsearch using Spring Data

To connect to Elasticsearch you must provide the address of one or more cluster nodes. The address can be specified by setting the `spring.data.elasticsearch.cluster-nodes` property to a comma-separated 'host:port' list. With this configuration in place, an `ElasticsearchTemplate` or `TransportClient` can be injected like any other Spring bean:

```
spring.data.elasticsearch.cluster-nodes=localhost:9300
```

```
@Component
public class MyBean {

    private final ElasticsearchTemplate template;

    public MyBean(ElasticsearchTemplate template) {
        this.template = template;
    }

    // ...
}
```

If you add your own `ElasticsearchTemplate` or `TransportClient` `@Bean` it will replace the default.

Spring Data Elasticsearch repositories

Spring Data includes repository support for Elasticsearch. As with the JPA repositories discussed earlier, the basic principle is that queries are constructed for you automatically based on method names.

In fact, both Spring Data JPA and Spring Data Elasticsearch share the same common infrastructure; so you could take the JPA example from earlier and, assuming that `City` is now an Elasticsearch `@Document` class rather than a JPA `@Entity`, it will work in the same way.

Tip

For complete details of Spring Data Elasticsearch, refer to their [reference documentation](#).

30.7 Cassandra

[Cassandra](#) is an open source, distributed database management system designed to handle large amounts of data across many commodity servers. Spring Boot offers auto-configuration for Cassandra

and abstractions on top of it provided by [Spring Data Cassandra](#). There is a `spring-boot-starter-data-cassandra` ‘Starter’ for collecting the dependencies in a convenient way.

Connecting to Cassandra

You can inject an auto-configured `CassandraTemplate` or a `Cassandra Session` instance as you would with any other Spring Bean. The `spring.data.cassandra.*` properties can be used to customize the connection. Generally you will provide `keyspace-name` and `contact-points` properties:

```
spring.data.cassandra.keyspace-name=mykeyspace
spring.data.cassandra.contact-points=cassandrahost1,cassandrahost2
```

```
@Component
public class MyBean {

    private CassandraTemplate template;

    @Autowired
    public MyBean(CassandraTemplate template) {
        this.template = template;
    }

    // ...
}
```

If you add a `@Bean` of your own of type `CassandraTemplate` it will replace the default.

Spring Data Cassandra repositories

Spring Data includes basic repository support for Cassandra. Currently this is more limited than the JPA repositories discussed earlier, and will need to annotate finder methods with `@Query`.

Tip

For complete details of Spring Data Cassandra, refer to their [reference documentation](#).

30.8 Couchbase

[Couchbase](#) is an open-source, distributed multi-model NoSQL document-oriented database that is optimized for interactive applications. Spring Boot offers auto-configuration for Couchbase and abstractions on top of it provided by [Spring Data Couchbase](#). There is a `spring-boot-starter-data-couchbase` ‘Starter’ for collecting the dependencies in a convenient way.

Connecting to Couchbase

You can very easily get a `Bucket` and `Cluster` by adding the Couchbase SDK and some configuration. The `spring.couchbase.*` properties can be used to customize the connection. Generally you will provide the bootstrap hosts, bucket name and password:

```
spring.couchbase.bootstrap-hosts=my-host-1,192.168.1.123
spring.couchbase.bucket.name=my-bucket
spring.couchbase.bucket.password=secret
```


Tip

You need to provide *at least* the bootstrap host(s), in which case the bucket name is default and the password is the empty String. Alternatively, you can define your own `org.springframework.data.couchbase.config.CouchbaseConfigurer` `@Bean` to take control over the whole configuration.

It is also possible to customize some of the `CouchbaseEnvironment` settings. For instance the following configuration changes the timeout to use to open a new `Bucket` and enables SSL support:

```
spring.couchbase.env.timeouts.connect=3000
spring.couchbase.env.ssl.key-store=/location/of/keystore.jks
spring.couchbase.env.ssl.key-store-password=secret
```

Check the `spring.couchbase.env.*` properties for more details.

Spring Data Couchbase repositories

Spring Data includes repository support for Couchbase. For complete details of Spring Data Couchbase, refer to their [reference documentation](#).

You can inject an auto-configured `CouchbaseTemplate` instance as you would with any other Spring Bean as long as a *default* `CouchbaseConfigurer` is available (that happens when you enable the couchbase support as explained above).

```
@Component
public class MyBean {

    private final CouchbaseTemplate template;

    @Autowired
    public MyBean(CouchbaseTemplate template) {
        this.template = template;
    }

    // ...

}
```

There are a few beans that you can define in your own configuration to override those provided by the auto-configuration:

- A `CouchbaseTemplate` `@Bean` with name `couchbaseTemplate`
- An `IndexManager` `@Bean` with name `couchbaseIndexManager`
- A `CustomConversions` `@Bean` with name `couchbaseCustomConversions`

To avoid hard-coding those names in your own config, you can reuse `BeanNames` provided by Spring Data Couchbase. For instance, you can customize the converters to use as follows:

```
@Configuration
public class SomeConfiguration {

    @Bean(BeanNames.COUCBASE_CUSTOM_CONVERSIONS)
    public CustomConversions myCustomConversions() {
        return new CustomConversions(...);
    }

    // ...

}
```

Tip

If you want to fully bypass the auto-configuration for Spring Data Couchbase, provide your own `org.springframework.data.couchbase.config.AbstractCouchbaseDataConfiguration` implementation.

30.9 LDAP

[LDAP](#) (Lightweight Directory Access Protocol) is an open, vendor-neutral, industry standard application protocol for accessing and maintaining distributed directory information services over an IP network. Spring Boot offers auto-configuration for any compliant LDAP server as well as support for the embedded in-memory LDAP server from [UnboundID](#).

LDAP abstractions are provided by [Spring Data LDAP](#). There is a `spring-boot-starter-data-ldap` ‘Starter’ for collecting the dependencies in a convenient way.

Connecting to an LDAP server

To connect to an LDAP server make sure you declare a dependency on the `spring-boot-starter-data-ldap` ‘Starter’ or `spring-ldap-core` then declare the URLs of your server in your `application.properties`:

```
spring.ldap.urls=ldap://myserver:1235
spring.ldap.username=admin
spring.ldap.password=secret
```

If you need to customize connection settings you can use the `spring.ldap.base` and `spring.ldap.base-environment` properties.

Spring Data LDAP repositories

Spring Data includes repository support for LDAP. For complete details of Spring Data LDAP, refer to their [reference documentation](#).

You can also inject an auto-configured `LdapTemplate` instance as you would with any other Spring Bean.

```
@Component
public class MyBean {

    private final LdapTemplate template;

    @Autowired
    public MyBean(LdapTemplate template) {
        this.template = template;
    }

    // ...
}
```

Embedded in-memory LDAP server

For testing purposes Spring Boot supports auto-configuration of an in-memory LDAP server from [UnboundID](#). To configure the server add a dependency to `com.unboundid:unboundid-ldapsdk` and declare a `base-dn` property:

```
spring.ldap.embedded.base-dn=dc=spring,dc=io
```

By default the server will start on a random port and they trigger the regular LDAP support (there is no need to specify a `spring.ldap.urls` property).

If there is a `schema.ldif` file on your classpath it will be used to initialize the server. You can also use the `spring.ldap.embedded.ldif` property if you want to load the initialization script from a different resource.

By default, a standard schema will be used to validate LDIF files, you can turn off validation altogether using the `spring.ldap.embedded.validation.enabled` property. If you have custom attributes, you can use `spring.ldap.embedded.validation.schema` to define your custom attribute types or object classes.

30.10 InfluxDB

[InfluxDB](#) is an open-source time series database optimized for fast, high-availability storage and retrieval of time series data in fields such as operations monitoring, application metrics, Internet of Things sensor data, and real-time analytics.

Connecting to InfluxDB

Spring Boot auto-configures an InfluxDB instance as long as the `influxdb-java` client is on the classpath and the url of the database is set:

```
spring.influx.url=http://172.0.0.1:8086
```

If the connection to InfluxDB requires a user and password, you can set the `spring.influx.user` and `spring.influx.password` properties accordingly.

InfluxDB relies on OkHttp. If you need to tune the http client InfluxDB uses behind the scenes, you can register a `OkHttpClient.Builder` bean.

31. Caching

The Spring Framework provides support for transparently adding caching to an application. At its core, the abstraction applies caching to methods, reducing thus the number of executions based on the information available in the cache. The caching logic is applied transparently, without any interference to the invoker. Spring Boot auto-configures the cache infrastructure as long as the caching support is enabled via the `@EnableCaching` annotation.

Note

Check the [relevant section](#) of the Spring Framework reference for more details.

In a nutshell, adding caching to an operation of your service is as easy as adding the relevant annotation to its method:

```
import org.springframework.cache.annotation.Cacheable
import org.springframework.stereotype.Component;

@Component
public class MathService {

    @Cacheable("piDecimals")
    public int computePiDecimal(int i) {
        // ...
    }
}
```

This example demonstrates the use of caching on a potentially costly operation. Before invoking `computePiDecimal`, the abstraction will look for an entry in the `piDecimals` cache matching the `i` argument. If an entry is found, the content in the cache is immediately returned to the caller and the method is not invoked. Otherwise, the method is invoked and the cache is updated before returning the value.

Note

You can also use the standard JSR-107 (JCache) annotations (e.g. `@CacheResult`) transparently. We strongly advise you however to not mix and match them.

If you do not add any specific cache library, Spring Boot will auto-configure a [Simple provider](#) that uses concurrent maps in memory. When a cache is required (i.e. `piDecimals` in the example above), this provider will create it on-the-fly for you. The simple provider is not really recommended for production usage, but it's great for getting started and making sure that you understand the features. When you have made up your mind about the cache provider to use, please make sure to read its documentation to figure out how to configure the caches that your application uses. Practically all providers require you to explicitly configure every cache that you use in the application. Some offer a way to customize the default caches defined by the `spring.cache.cache-names` property.

Tip

It is also possible to [update](#) or [evict](#) data from the cache transparently.

Note

If you are using the cache infrastructure with beans that are not interface-based, make sure to enable the `proxyTargetClass` attribute of `@EnableCaching`.

31.1 Supported cache providers

The cache abstraction does not provide an actual store and relies on abstraction materialized by the `org.springframework.cache.Cache` and `org.springframework.cache.CacheManager` interfaces.

If you haven't defined a bean of type `CacheManager` or a `CacheResolver` named `cacheResolver` (see `CachingConfigurer`), Spring Boot tries to detect the following providers (in this order):

- [Generic](#)
- [JCache \(JSR-107\)](#) (EhCache 3, Hazelcast, Infinispan, etc)
- [EhCache 2.x](#)
- [Hazelcast](#)
- [Infinispan](#)
- [Couchbase](#)
- [Redis](#)
- [Caffeine](#)
- [Simple](#)

Tip

It is also possible to *force* the cache provider to use via the `spring.cache.type` property. Use this property if you need to [disable caching altogether](#) in certain environment (e.g. tests).

Tip

Use the `spring-boot-starter-cache` 'Starter' to quickly add basic caching dependencies. The starter brings in `spring-context-support`: if you are adding dependencies manually, you must include `spring-context-support` in order to use the JCache, EhCache 2.x or Guava support.

If the `CacheManager` is auto-configured by Spring Boot, you can further tune its configuration before it is fully initialized by exposing a bean implementing the `CacheManagerCustomizer` interface. The following sets a flag to say that null values should be passed down to the underlying map.

```
@Bean
public CacheManagerCustomizer<ConcurrentMapCacheManager> cacheManagerCustomizer() {
    return new CacheManagerCustomizer<ConcurrentMapCacheManager>() {
        @Override
        public void customize(ConcurrentMapCacheManager cacheManager) {
            cacheManager.setAllowNullValues(false);
        }
    };
}
```

Note

In the example above, an auto-configured `ConcurrentMapCacheManager` is expected. If that is not the case (either you provided your own config or a different cache provider was auto-configured), the customizer won't be invoked at all. You can have as many customizers as you want and you can also order them as usual using `@Order` or `Ordered`.

Generic

Generic caching is used if the context defines *at least one* `org.springframework.cache.Cache` bean. A `CacheManager` wrapping all beans of that type is created.

JCache (JSR-107)

JCache is bootstrapped via the presence of a `javax.cache.spi.CachingProvider` on the classpath (i.e. a JSR-107 compliant caching library) and the `JCacheCacheManager` provided by the `spring-boot-starter-cache` 'Starter'. There are various compliant libraries out there and Spring Boot provides dependency management for Ehcache 3, Hazelcast and Infinispan. Any other compliant library can be added as well.

It might happen that more than one provider is present, in which case the provider must be explicitly specified. Even if the JSR-107 standard does not enforce a standardized way to define the location of the configuration file, Spring Boot does its best to accommodate with implementation details.

```
# Only necessary if more than one provider is present
spring.cache.jcache.provider=com.acme.MyCachingProvider
spring.cache.jcache.config=classpath:acme.xml
```

Note

Since a cache library may offer both a native implementation and JSR-107 support Spring Boot will prefer the JSR-107 support so that the same features are available if you switch to a different JSR-107 implementation.

Tip

Spring Boot has a [general support for Hazelcast](#). If a single `HazelcastInstance` is available, it is automatically reused for the `CacheManager` as well unless the `spring.cache.jcache.config` property is specified.

There are several ways to customize the underlying `javax.cache.cacheManager`:

- Caches can be created on startup via the `spring.cache.cache-names` property. If a custom `javax.cache.configuration.Configuration` bean is defined, it is used to customize them.
- `org.springframework.boot.autoconfigure.cache.JCacheManagerCustomizer` beans are invoked with the reference of the `CacheManager` for full customization.

Tip

If a standard `javax.cache.CacheManager` bean is defined, it is wrapped automatically in a `org.springframework.cache.CacheManager` implementation that the abstraction expects. No further customization is applied on it.

EhCache 2.x

EhCache 2.x is used if a file named `ehcache.xml` can be found at the root of the classpath. If EhCache 2.x, the `EhCacheCacheManager` provided by the `spring-boot-starter-cache` 'Starter' and such file is present it is used to bootstrap the cache manager. An alternate configuration file can be provided as well using:

```
spring.cache.ehcache.config=classpath:config/another-config.xml
```

Hazelcast

Spring Boot has a [general support for Hazelcast](#). If a `HazelcastInstance` has been auto-configured, it is automatically wrapped in a `CacheManager`.

Infinispan

Infinispan has no default configuration file location so it must be specified explicitly (or the default bootstrap is used).

```
spring.cache.infinispan.config=infinispan.xml
```

Caches can be created on startup via the `spring.cache.cache-names` property. If a custom `ConfigurationBuilder` bean is defined, it is used to customize them.

Note

The support of Infinispan in Spring Boot is restricted to the embedded mode and is quite basic. If you want more options you should use the official Infinispan Spring Boot starter instead, check [the documentation](#) for more details.

Couchbase

If the Couchbase java client and the `couchbase-spring-cache` implementation are available and Couchbase is [configured](#), a `CouchbaseCacheManager` will be auto-configured. It is also possible to create additional caches on startup using the `spring.cache.cache-names` property. These will operate on the `Bucket` that was auto-configured. You can *also* create additional caches on another `Bucket` using the customizer: assume you need two caches on the "main" `Bucket` (`foo` and `bar`) and one `biz` cache with a custom time to live of 2sec on the another `Bucket`. First, you can create the two first caches simply via configuration:

```
spring.cache.cache-names=foo,bar
```

Then define this extra `@Configuration` to configure the extra `Bucket` and the `biz` cache:

```
@Configuration
public class CouchbaseCacheConfiguration {

    private final Cluster cluster;

    public CouchbaseCacheConfiguration(Cluster cluster) {
        this.cluster = cluster;
    }

    @Bean
    public Bucket anotherBucket() {
        return this.cluster.openBucket("another", "secret");
    }
}
```

```

    }

    @Bean
    public CacheManagerCustomizer<CouchbaseCacheManager> cacheManagerCustomizer() {
        return c -> {
            c.prepareCache("biz", CacheBuilder.newInstance(anotherBucket())
                .withExpiration(2));
        };
    }
}

```

This sample configuration reuses the `Cluster` that was created via auto-configuration.

Redis

If Redis is available and configured, the `RedisCacheManager` is auto-configured. It is also possible to create additional caches on startup using the `spring.cache.cache-names` property.

Note

By default, a key prefix is added to prevent that if two separate caches use the same key, Redis would have overlapping keys and be likely to return invalid values. We strongly recommend to keep this setting enabled if you create your own `RedisCacheManager`.

Caffeine

Caffeine is a Java 8 rewrite of Guava's cache that supersedes the Guava support. If Caffeine is present, a `CaffeineCacheManager` (provided by the `spring-boot-starter-cache` 'Starter') is auto-configured. Caches can be created on startup using the `spring.cache.cache-names` property and customized by one of the following (in this order):

1. A cache spec defined by `spring.cache.caffeine.spec`
2. A `com.github.benmanes.caffeine.cache.CaffeineSpec` bean is defined
3. A `com.github.benmanes.caffeine.cache.Caffeine` bean is defined

For instance, the following configuration creates a `foo` and `bar` caches with a maximum size of 500 and a *time to live* of 10 minutes

```

spring.cache.cache-names=foo,bar
spring.cache.caffeine.spec=maximumSize=500,expireAfterAccess=600s

```

Besides, if a `com.github.benmanes.caffeine.cache.CacheLoader` bean is defined, it is automatically associated to the `CaffeineCacheManager`. Since the `CacheLoader` is going to be associated to *all* caches managed by the cache manager, it must be defined as `CacheLoader<Object, Object>`. Any other generic type will be ignored by the auto-configuration.

Simple

If none of the other providers can be found, a simple implementation using a `ConcurrentHashMap` as cache store is configured. This is the default if no caching library is present in your application. Caches are created on-the-fly by default but you can restrict the list of available caches using the `cache-names` property. For instance, if you want only `foo` and `bar` caches:

```

spring.cache.cache-names=foo,bar

```


If you do this and your application uses a cache not listed then it will fail at runtime when the cache is needed, but not on startup. This is similar to the way the "real" cache providers behave if you use an undeclared cache.

None

When `@EnableCaching` is present in your configuration, a suitable cache configuration is expected as well. If you need to disable caching altogether in certain environments, force the cache type to `none` to use a no-op implementation:

```
spring.cache.type=none
```

32. Messaging

The Spring Framework provides extensive support for integrating with messaging systems: from simplified use of the JMS API using `JmsTemplate` to a complete infrastructure to receive messages asynchronously. Spring AMQP provides a similar feature set for the ‘Advanced Message Queuing Protocol’ and Spring Boot also provides auto-configuration options for `RabbitTemplate` and RabbitMQ. There is also support for STOMP messaging natively in Spring WebSocket and Spring Boot has support for that through starters and a small amount of auto-configuration. Spring Boot also has support for Apache Kafka.

32.1 JMS

The `javax.jms.ConnectionFactory` interface provides a standard method of creating a `javax.jms.Connection` for interacting with a JMS broker. Although Spring needs a `ConnectionFactory` to work with JMS, you generally won’t need to use it directly yourself and you can instead rely on higher level messaging abstractions (see the [relevant section](#) of the Spring Framework reference documentation for details). Spring Boot also auto-configures the necessary infrastructure to send and receive messages.

ActiveMQ support

Spring Boot can also configure a `ConnectionFactory` when it detects that ActiveMQ is available on the classpath. If the broker is present, an embedded broker is started and configured automatically (as long as no broker URL is specified through configuration).

Note

If you are using `spring-boot-starter-activemq` the necessary dependencies to connect or embed an ActiveMQ instance are provided, as well as the Spring infrastructure to integrate with JMS.

ActiveMQ configuration is controlled by external configuration properties in `spring.activemq.*`. For example, you might declare the following section in `application.properties`:

```
spring.activemq.broker-url=tcp://192.168.1.210:9876
spring.activemq.user=admin
spring.activemq.password=secret
```

You can also pool JMS resources by adding a dependency to `org.apache.activemq:activemq-pool` and configure the `PooledConnectionFactory` accordingly:

```
spring.activemq.pool.enabled=true
spring.activemq.pool.max-connections=50
```

Tip

See [ActiveMQProperties](#) for more of the supported options. You can also register an arbitrary number of beans implementing `ActiveMQConnectionFactoryCustomizer` for more advanced customizations.

By default, ActiveMQ creates a destination if it does not exist yet, so destinations are resolved against their provided names.

Artemis support

Spring Boot can auto-configure a `ConnectionFactory` when it detects that Artemis is available on the classpath. If the broker is present, an embedded broker is started and configured automatically (unless the mode property has been explicitly set). The supported modes are: `embedded` (to make explicit that an embedded broker is required and should lead to an error if the broker is not available in the classpath), and `native` to connect to a broker using the `netty` transport protocol. When the latter is configured, Spring Boot configures a `ConnectionFactory` connecting to a broker running on the local machine with the default settings.

Note

If you are using `spring-boot-starter-artemis` the necessary dependencies to connect to an existing Artemis instance are provided, as well as the Spring infrastructure to integrate with JMS. Adding `org.apache.activemq:artemis-jms-server` to your application allows you to use the embedded mode.

Artemis configuration is controlled by external configuration properties in `spring.artemis.*`. For example, you might declare the following section in `application.properties`:

```
spring.artemis.mode=native
spring.artemis.host=192.168.1.210
spring.artemis.port=9876
spring.artemis.user=admin
spring.artemis.password=secret
```

When embedding the broker, you can choose if you want to enable persistence, and the list of destinations that should be made available. These can be specified as a comma-separated list to create them with the default options; or you can define bean(s) of type `org.apache.activemq.artemis.jms.server.config.JMSQueueConfiguration` or `org.apache.activemq.artemis.jms.server.config.TopicConfiguration`, for advanced queue and topic configurations respectively.

See [ArtemisProperties](#) for more of the supported options.

No JNDI lookup is involved at all and destinations are resolved against their names, either using the 'name' attribute in the Artemis configuration or the names provided through configuration.

Using a JNDI ConnectionFactory

If you are running your application in an Application Server Spring Boot will attempt to locate a JMS `ConnectionFactory` using JNDI. By default the locations `java:/JmsXA` and `java:/XAConnectionFactory` will be checked. You can use the `spring.jms.jndi-name` property if you need to specify an alternative location:

```
spring.jms.jndi-name=java:/MyConnectionFactory
```

Sending a message

Spring's `JmsTemplate` is auto-configured and you can autowire it directly into your own beans:

```
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.jms.core.JmsTemplate;
import org.springframework.stereotype.Component;
```

```

@Component
public class MyBean {

    private final JmsTemplate jmsTemplate;

    @Autowired
    public MyBean(JmsTemplate jmsTemplate) {
        this.jmsTemplate = jmsTemplate;
    }

    // ...

}

```

Note

[JmsMessagingTemplate](#) can be injected in a similar manner. If a `DestinationResolver` or `MessageConverter` beans are defined, they are associated automatically to the auto-configured `JmsTemplate`.

Receiving a message

When the JMS infrastructure is present, any bean can be annotated with `@JmsListener` to create a listener endpoint. If no `JmsListenerContainerFactory` has been defined, a default one is configured automatically. If a `DestinationResolver` or `MessageConverter` beans are defined, they are associated automatically to the default factory.

The default factory is transactional by default. If you are running in an infrastructure where a `JtaTransactionManager` is present, it will be associated to the listener container by default. If not, the `sessionTransacted` flag will be enabled. In that latter scenario, you can associate your local data store transaction to the processing of an incoming message by adding `@Transactional` on your listener method (or a delegate thereof). This will make sure that the incoming message is acknowledged once the local transaction has completed. This also includes sending response messages that have been performed on the same JMS session.

The following component creates a listener endpoint on the `someQueue` destination:

```

@Component
public class MyBean {

    @JmsListener(destination = "someQueue")
    public void processMessage(String content) {
        // ...
    }

}

```

Tip

Check [the Javadoc of @EnableJms](#) for more details.

If you need to create more `JmsListenerContainerFactory` instances or if you want to override the default, Spring Boot provides a `DefaultJmsListenerContainerFactoryConfigurer` that you can use to initialize a `DefaultJmsListenerContainerFactory` with the same settings as the one that is auto-configured.

For instance, the following exposes another factory that uses a specific `MessageConverter`:

```

@Configuration
static class JmsConfiguration {

    @Bean
    public DefaultJmsListenerContainerFactory myFactory(
        DefaultJmsListenerContainerFactoryConfigurer configurer) {
        DefaultJmsListenerContainerFactory factory =
            new DefaultJmsListenerContainerFactory();
        configurer.configure(factory, connectionFactory());
        factory.setMessageConverter(myMessageConverter());
        return factory;
    }
}

```

Then you can use in any `@JmsListener`-annotated method as follows:

```

@Component
public class MyBean {

    @JmsListener(destination = "someQueue", containerFactory="myFactory")
    public void processMessage(String content) {
        // ...
    }
}

```

32.2 AMQP

The Advanced Message Queuing Protocol (AMQP) is a platform-neutral, wire-level protocol for message-oriented middleware. The Spring AMQP project applies core Spring concepts to the development of AMQP-based messaging solutions. Spring Boot offers several conveniences for working with AMQP via RabbitMQ, including the `spring-boot-starter-amqp` ‘Starter’.

RabbitMQ support

RabbitMQ is a lightweight, reliable, scalable and portable message broker based on the AMQP protocol. Spring uses `RabbitMQ` to communicate using the AMQP protocol.

RabbitMQ configuration is controlled by external configuration properties in `spring.rabbitmq.*`. For example, you might declare the following section in `application.properties`:

```

spring.rabbitmq.host=localhost
spring.rabbitmq.port=5672
spring.rabbitmq.username=admin
spring.rabbitmq.password=secret

```

See [RabbitProperties](#) for more of the supported options.

Tip

Check [Understanding AMQP, the protocol used by RabbitMQ](#) for more details.

Sending a message

Spring's `AmqpTemplate` and `AmqpAdmin` are auto-configured and you can autowire them directly into your own beans:

```

import org.springframework.amqp.core.AmqpAdmin;

```

```

import org.springframework.amqp.core.AmqpTemplate;
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.stereotype.Component;

@Component
public class MyBean {

    private final AmqpAdmin amqpAdmin;
    private final AmqpTemplate amqpTemplate;

    @Autowired
    public MyBean(AmqpAdmin amqpAdmin, AmqpTemplate amqpTemplate) {
        this.amqpAdmin = amqpAdmin;
        this.amqpTemplate = amqpTemplate;
    }

    // ...
}

```

Note

[RabbitMessagingTemplate](#) can be injected in a similar manner. If a `MessageConverter` bean is defined, it is associated automatically to the auto-configured `AmqpTemplate`.

Any `org.springframework.amqp.core.Queue` that is defined as a bean will be automatically used to declare a corresponding queue on the RabbitMQ instance if necessary.

You can enable retries on the `AmqpTemplate` to retry operations, for example in the event the broker connection is lost. Retries are disabled by default.

Receiving a message

When the Rabbit infrastructure is present, any bean can be annotated with `@RabbitListener` to create a listener endpoint. If no `RabbitListenerContainerFactory` has been defined, a default `SimpleRabbitListenerContainerFactory` is configured automatically and you can switch to a direct container using the `spring.rabbitmq.listener.type` property. If a `MessageConverter` or `MessageRecoverer` beans are defined, they are associated automatically to the default factory.

The following component creates a listener endpoint on the `someQueue` queue:

```

@Component
public class MyBean {

    @RabbitListener(queues = "someQueue")
    public void processMessage(String content) {
        // ...
    }
}

```

Tip

Check [the Javadoc of @EnableRabbit](#) for more details.

If you need to create more `RabbitListenerContainerFactory` instances or if you want to override the default, Spring Boot provides a `SimpleRabbitListenerContainerFactoryConfigurer` and `DirectRabbitListenerContainerFactoryConfigurer` that you can use to initialize a `SimpleRabbitListenerContainerFactory` and

`DirectRabbitListenerContainerFactory` with the same settings as the one used by the auto-configuration.

Tip

It doesn't matter which container type you've chosen, those two beans are exposed by the auto-configuration.

For instance, the following exposes another factory that uses a specific `MessageConverter`:

```
@Configuration
static class RabbitConfiguration {

    @Bean
    public SimpleRabbitListenerContainerFactory myFactory(
        SimpleRabbitListenerContainerFactoryConfigurer configurer) {
        SimpleRabbitListenerContainerFactory factory =
            new SimpleRabbitListenerContainerFactory();
        configurer.configure(factory, connectionFactory);
        factory.setMessageConverter(myMessageConverter());
        return factory;
    }
}
```

Then you can use in any `@RabbitListener`-annotated method as follows:

```
@Component
public class MyBean {

    @RabbitListener(queues = "someQueue", containerFactory="myFactory")
    public void processMessage(String content) {
        // ...
    }
}
```

You can enable retries to handle situations where your listener throws an exception. By default `RejectAndDontRequeueRecoverer` is used but you can define a `MessageRecoverer` of your own. When retries are exhausted, the message will be rejected and either dropped or routed to a dead-letter exchange if the broker is configured so. Retries are disabled by default.

Important

If retries are not enabled and the listener throws an exception, by default the delivery will be retried indefinitely. You can modify this behavior in two ways; set the `defaultRequeueRejected` property to `false` and zero re-deliveries will be attempted; or, throw an `AmqpRejectAndDontRequeueException` to signal the message should be rejected. This is the mechanism used when retries are enabled and the maximum delivery attempts are reached.

32.3 Apache Kafka Support

[Apache Kafka](#) is supported by providing auto-configuration of the `spring-kafka` project.

Kafka configuration is controlled by external configuration properties in `spring.kafka.*`. For example, you might declare the following section in `application.properties`:

```
spring.kafka.bootstrap.servers=localhost:9092
spring.kafka.consumer.group-id=myGroup
```

See [KafkaProperties](#) for more of the supported options.

Sending a Message

Spring's `KafkaTemplate` is auto-configured and you can autowire them directly in your own beans:

```
@Component
public class MyBean {

    private final KafkaTemplate kafkaTemplate;

    @Autowired
    public MyBean(KafkaTemplate kafkaTemplate) {
        this.kafkaTemplate = kafkaTemplate;
    }

    // ...
}
```

Receiving a Message

When the Apache Kafka infrastructure is present, any bean can be annotated with `@KafkaListener` to create a listener endpoint. If no `KafkaListenerContainerFactory` has been defined, a default one is configured automatically with keys defined in `spring.kafka.listener.*`.

The following component creates a listener endpoint on the `someTopic` topic:

```
@Component
public class MyBean {

    @KafkaListener(topics = "someTopic")
    public void processMessage(String content) {
        // ...
    }

}
```

Additional Kafka Properties

The properties supported by auto configuration are shown in [Appendix A, Common application properties](#). Note that these properties (hyphenated or camelCase) map directly to the Apache Kafka dotted properties for the most part, refer to the Apache Kafka documentation for details.

The first few of these properties apply to both producers and consumers, but can be specified at the producer or consumer level if you wish to use different values for each. Apache Kafka designates properties with an importance: HIGH, MEDIUM and LOW. Spring Boot auto configuration supports all HIGH importance properties, some selected MEDIUM and LOW, and any that do not have a default value.

Only a subset of the properties supported by Kafka are available via the `KafkaProperties` class. If you wish to configure the producer or consumer with additional properties that are not directly supported, use the following:

```
spring.kafka.properties.foo.bar=baz
spring.kafka.consumer.properties.fiz.buz=qux
spring.kafka.producer.properties.baz.qux=fiz
```


This sets the common `foo.bar` Kafka property to `baz` (applies to both producers and consumers), the consumer `fiz.buz` property to `qux` and the `baz.qux` producer property to `fiz`.

Important

Properties set in this way will override any configuration item that Spring Boot explicitly supports.

33. Calling REST services with ‘RestTemplate’

If you need to call remote REST services from your application, you can use Spring Framework’s `RestTemplate` class. Since `RestTemplate` instances often need to be customized before being used, Spring Boot does not provide any single auto-configured `RestTemplate` bean. It does, however, auto-configure a `RestTemplateBuilder` which can be used to create `RestTemplate` instances when needed. The auto-configured `RestTemplateBuilder` will ensure that sensible `HttpMessageConverters` are applied to `RestTemplate` instances.

Here’s a typical example:

```
@Service
public class MyService {

    private final RestTemplate restTemplate;

    public MyBean(RestTemplateBuilder restTemplateBuilder) {
        this.restTemplate = restTemplateBuilder.build();
    }

    public Details someRestCall(String name) {
        return this.restTemplate.getForObject("/{name}/details", Details.class, name);
    }
}
```

Tip

`RestTemplateBuilder` includes a number of useful methods that can be used to quickly configure a `RestTemplate`. For example, to add BASIC auth support you can use `builder.basicAuthorization("user", "password").build()`.

33.1 RestTemplate customization

There are three main approaches to `RestTemplate` customization, depending on how broadly you want the customizations to apply.

To make the scope of any customizations as narrow as possible, inject the auto-configured `RestTemplateBuilder` and then call its methods as required. Each method call returns a new `RestTemplateBuilder` instance so the customizations will only affect this use of the builder.

To make an application-wide, additive customization a `RestTemplateCustomizer` bean can be used. All such beans are automatically registered with the auto-configured `RestTemplateBuilder` and will be applied to any templates that are built with it.

Here’s an example of a customizer that configures the use of a proxy for all hosts except `192.168.0.5`:

```
static class ProxyCustomizer implements RestTemplateCustomizer {

    @Override
    public void customize(RestTemplate restTemplate) {
        HttpHost proxy = new HttpHost("proxy.example.com");
        HttpClient httpClient = HttpClientBuilder.create()
            .setRoutePlanner(new DefaultProxyRoutePlanner(proxy) {

                @Override
                public HttpHost determineProxy(HttpHost target,
                    HttpRequest request, HttpContext context)
            }
        );
    }
}
```

```
        throws HttpException {
            if (target.getHostName().equals("192.168.0.5")) {
                return null;
            }
            return super.determineProxy(target, request, context);
        }
    }

    }).build();
    restTemplate.setRequestFactory(
        new HttpComponentsClientHttpRequestFactory(httpClient));
}
}
```

Lastly, the most extreme (and rarely used) option is to create your own `RestTemplateBuilder` bean. This will switch off the auto-configuration of a `RestTemplateBuilder` and will prevent any `RestTemplateCustomizer` beans from being used.

34. Calling REST services with ‘WebClient’

If you have Spring WebFlux on your classpath, you can also choose to use `WebClient` to call remote REST services; compared to `RestTemplate`, this client has more a functional feel to it and is fully reactive. You can create your own client instance with the builder `WebClient.create()`, which already provides a good out-of-the-box experience. See the [relevant section on WebClient](#).

Spring Boot will create and pre-configure such a builder for you; for example, client HTTP codecs will be configured just like the server ones (see [WebFlux HTTP codecs auto-configuration](#)).

Here’s a typical example:

```
@Service
public class MyService {

    private final WebClient webClient;

    public MyBean(WebClient.Builder webClientBuilder) {
        this.webClient = webClientBuilder.baseUrl("http://example.org").build();
    }

    public Mono<Details> someRestCall(String name) {
        return this.webClient.get().url("/{name}/details", name)
            .retrieve().bodyToMono(Details.class);
    }
}
```

34.1 WebClient customization

There are three main approaches to `WebClient` customization, depending on how broadly you want the customizations to apply.

To make the scope of any customizations as narrow as possible, inject the auto-configured `WebClient.Builder` and then call its methods as required. `WebClient.Builder` instances are stateful; any change on the builder will be reflected in all clients subsequently created with it. If you’d like to create several clients with the same builder, you can also consider cloning the builder with `WebClient.Builder other = builder.clone();`

To make an application-wide, additive customization to all `WebClient.Builder` instances, you can declare `WebClientCustomizer` beans and change the `WebClient.Builder` as you would do locally at the point of injection.

Lastly, you can fall back to the original API and just use `WebClient.create()`. In that case, no auto-configuration nor `WebClientCustomizer` will be applied.

35. Validation

The method validation feature supported by Bean Validation 1.1 is automatically enabled as long as a JSR-303 implementation (e.g. Hibernate validator) is on the classpath. This allows bean methods to be annotated with `javax.validation` constraints on their parameters and/or on their return value. Target classes with such annotated methods need to be annotated with the `@Validated` annotation at the type level for their methods to be searched for inline constraint annotations.

For instance, the following service triggers the validation of the first argument, making sure its size is between 8 and 10

```
@Service
@Validated
public class MyBean {

    public Archive findByCodeAndAuthor(@Size(min = 8, max = 10) String code,
        Author author) {
        ...
    }
}
```

36. Sending email

The Spring Framework provides an easy abstraction for sending email using the `JavaMailSender` interface and Spring Boot provides auto-configuration for it as well as a starter module.

Tip

Check the [reference documentation](#) for a detailed explanation of how you can use `JavaMailSender`.

If `spring.mail.host` and the relevant libraries (as defined by `spring-boot-starter-mail`) are available, a default `JavaMailSender` is created if none exists. The sender can be further customized by configuration items from the `spring.mail` namespace, see the [MailProperties](#) for more details.

In particular, certain default timeout values are infinite and you may want to change that to avoid having a thread blocked by an unresponsive mail server:

```
spring.mail.properties.mail.smtp.connectiontimeout=5000
spring.mail.properties.mail.smtp.timeout=3000
spring.mail.properties.mail.smtp.writetimeout=5000
```

37. Distributed Transactions with JTA

Spring Boot supports distributed JTA transactions across multiple XA resources using either an [Atomikos](#) or [Bitronix](#) embedded transaction manager. JTA transactions are also supported when deploying to a suitable Java EE Application Server.

When a JTA environment is detected, Spring's `JtaTransactionManager` will be used to manage transactions. Auto-configured JMS, DataSource and JPA beans will be upgraded to support XA transactions. You can use standard Spring idioms such as `@Transactional` to participate in a distributed transaction. If you are within a JTA environment and still want to use local transactions you can set the `spring.jta.enabled` property to `false` to disable the JTA auto-configuration.

37.1 Using an Atomikos transaction manager

Atomikos is a popular open source transaction manager which can be embedded into your Spring Boot application. You can use the `spring-boot-starter-jta-atomikos` Starter to pull in the appropriate Atomikos libraries. Spring Boot will auto-configure Atomikos and ensure that appropriate `depends-on` settings are applied to your Spring beans for correct startup and shutdown ordering.

By default Atomikos transaction logs will be written to a `transaction-logs` directory in your application home directory (the directory in which your application jar file resides). You can customize this directory by setting a `spring.jta.log-dir` property in your `application.properties` file. Properties starting `spring.jta.atomikos.properties` can also be used to customize the Atomikos `UserTransactionServiceImp`. See the [AtomikosProperties Javadoc](#) for complete details.

Note

To ensure that multiple transaction managers can safely coordinate the same resource managers, each Atomikos instance must be configured with a unique ID. By default this ID is the IP address of the machine on which Atomikos is running. To ensure uniqueness in production, you should configure the `spring.jta.transaction-manager-id` property with a different value for each instance of your application.

37.2 Using a Bitronix transaction manager

Bitronix is popular open source JTA transaction manager implementation. You can use the `spring-boot-starter-jta-bitronix` starter to add the appropriate Bitronix dependencies to your project. As with Atomikos, Spring Boot will automatically configure Bitronix and post-process your beans to ensure that startup and shutdown ordering is correct.

By default Bitronix transaction log files (`part1.btm` and `part2.btm`) will be written to a `transaction-logs` directory in your application home directory. You can customize this directory by using the `spring.jta.log-dir` property. Properties starting `spring.jta.bitronix.properties` are also bound to the `bitronix.tm.Configuration` bean, allowing for complete customization. See the [Bitronix documentation](#) for details.

Note

To ensure that multiple transaction managers can safely coordinate the same resource managers, each Bitronix instance must be configured with a unique ID. By default this ID is the IP address

of the machine on which Bitronix is running. To ensure uniqueness in production, you should configure the `spring.jta.transaction-manager-id` property with a different value for each instance of your application.

37.3 Using a Narayana transaction manager

Narayana is popular open source JTA transaction manager implementation supported by JBoss. You can use the `spring-boot-starter-jta-narayana` starter to add the appropriate Narayana dependencies to your project. As with Atomikos and Bitronix, Spring Boot will automatically configure Narayana and post-process your beans to ensure that startup and shutdown ordering is correct.

By default Narayana transaction logs will be written to a `transaction-logs` directory in your application home directory (the directory in which your application jar file resides). You can customize this directory by setting a `spring.jta.log-dir` property in your `application.properties` file. Properties starting `spring.jta.narayana.properties` can also be used to customize the Narayana configuration. See the [NarayanaProperties Javadoc](#) for complete details.

Note

To ensure that multiple transaction managers can safely coordinate the same resource managers, each Narayana instance must be configured with a unique ID. By default this ID is set to 1. To ensure uniqueness in production, you should configure the `spring.jta.transaction-manager-id` property with a different value for each instance of your application.

37.4 Using a Java EE managed transaction manager

If you are packaging your Spring Boot application as a `war` or `ear` file and deploying it to a Java EE application server, you can use your application servers built-in transaction manager. Spring Boot will attempt to auto-configure a transaction manager by looking at common JNDI locations (`java:comp/UserTransaction`, `java:comp/TransactionManager` etc). If you are using a transaction service provided by your application server, you will generally also want to ensure that all resources are managed by the server and exposed over JNDI. Spring Boot will attempt to auto-configure JMS by looking for a `ConnectionFactory` at the JNDI path `java:/JmsXA` or `java:/XAConnectionFactory` and you can use the [spring.datasource.jndi-name property](#) to configure your `DataSource`.

37.5 Mixing XA and non-XA JMS connections

When using JTA, the primary JMS `ConnectionFactory` bean will be XA aware and participate in distributed transactions. In some situations you might want to process certain JMS messages using a non-XA `ConnectionFactory`. For example, your JMS processing logic might take longer than the XA timeout.

If you want to use a non-XA `ConnectionFactory` you can inject the `nonXaJmsConnectionFactory` bean rather than the `@Primary jmsConnectionFactory` bean. For consistency the `jmsConnectionFactory` bean is also provided using the bean alias `xaJmsConnectionFactory`.

For example:

```
// Inject the primary (XA aware) ConnectionFactory
@Autowired
private ConnectionFactory defaultConnectionFactory;
```



```
// Inject the XA aware ConnectionFactory (uses the alias and injects the same as above)
@Autowired
@Qualifier("xaJmsConnectionFactory")
private ConnectionFactory xaConnectionFactory;

// Inject the non-XA aware ConnectionFactory
@Autowired
@Qualifier("nonXaJmsConnectionFactory")
private ConnectionFactory nonXaConnectionFactory;
```

37.6 Supporting an alternative embedded transaction manager

The [XAConnectionFactoryWrapper](#) and [XADataSourceWrapper](#) interfaces can be used to support alternative embedded transaction managers. The interfaces are responsible for wrapping `XAConnectionFactory` and `XADataSource` beans and exposing them as regular `ConnectionFactory` and `DataSource` beans which will transparently enroll in the distributed transaction. `DataSource` and JMS auto-configuration will use JTA variants as long as you have a `JtaTransactionManager` bean and appropriate XA wrapper beans registered within your `ApplicationContext`.

The [BitronixXAConnectionFactoryWrapper](#) and [BitronixXADataSourceWrapper](#) provide good examples of how to write XA wrappers.

38. Hazelcast

If Hazelcast is on the classpath and a suitable configuration is found, Spring Boot will auto-configure a `HazelcastInstance` that you can inject in your application.

You can define a `com.hazelcast.config.Config` bean and we'll use that. If your configuration defines an instance name, we'll try to locate an existing instance rather than creating a new one.

You could also specify the `hazelcast.xml` configuration file to use via configuration:

```
spring.hazelcast.config=classpath:config/my-hazelcast.xml
```

Otherwise, Spring Boot tries to find the Hazelcast configuration from the default locations, that is `hazelcast.xml` in the working directory or at the root of the classpath. We also check if the `hazelcast.config` system property is set. Check the [Hazelcast documentation](#) for more details.

If `hazelcast-client` is present on the classpath, Spring Boot will first attempt to create a client with similar rules as above, that is:

- The presence of a `com.hazelcast.client.config.ClientConfig` bean
- A configuration file defined by the `spring.hazelcast.config` property
- The presence of the `hazelcast.client.config` system property
- A `hazelcast-client.xml` in the working directory or at the root of the classpath

Note

Spring Boot also has an [explicit caching support for Hazelcast](#). The `HazelcastInstance` is automatically wrapped in a `CacheManager` implementation if caching is enabled.

39. Quartz Scheduler

Spring Boot offers several conveniences for working with the Quartz scheduler, including the `spring-boot-starter-quartz` ‘Starter’. If Quartz is available, a `Scheduler` will be auto-configured (via the `SchedulerFactoryBean` abstraction).

Beans of the following types will be automatically picked up and associated with the the `Scheduler`:

- `JobDetail`: defines a particular `Job`. `JobDetail` instance can easily be built with the `JobBuilder` API
- `Calendar`
- `Trigger`: defines when a particular job is triggered

By default, an in-memory `JobStore` will be used. However, it is possible to configure a JDBC-based store if a `DataSource` bean is available in your application and if the `spring.quartz.job-store-type` property is configured accordingly:

```
spring.quartz.job-store-type=jdbc
```

When the jdbc store is used, the schema can be initialized on startup:

```
spring.quartz.jdbc.initialize-schema=true
```

Note

The database is detected by default and initialized using the standard scripts provided with the Quartz library. It is also possible to provide a custom script using the `spring.quartz.jdbc.schema` property.

Quartz Scheduler configuration can be customized using Quartz configuration properties (see `spring.quartz.properties.*`) and `SchedulerFactoryBeanCustomizer` beans which allow programmatic `SchedulerFactoryBean` customization.

Job can define setters to inject data map properties. Regular beans can also be injected in a similar manner:

```
public class SampleJob extends QuartzJobBean {

    private MyService myService;
    private String name;

    // Inject "MyService" bean
    public void setMyService(MyService myService) { ... }

    // Inject the "name" job data property
    public void setName(String name) { ... }

    @Override
    protected void executeInternal(JobExecutionContext context)
        throws JobExecutionException {
        ...
    }
}
```

40. Spring Integration

Spring Boot offers several conveniences for working with Spring Integration, including the `spring-boot-starter-integration` ‘Starter’. Spring Integration provides abstractions over messaging and also other transports such as HTTP, TCP etc. If Spring Integration is available on your classpath it will be initialized through the `@EnableIntegration` annotation.

Spring Boot will also configure some features that are triggered by the presence of additional Spring Integration modules. Message processing statistics will be published over JMX if `'spring-integration-jmx'` is also on the classpath. If `'spring-integration-jdbc'` is available, the default database schema can be created on startup:

```
spring.integration.jdbc.initializer.enabled=true
```

See the [IntegrationAutoConfiguration](#) and [IntegrationProperties](#) classes for more details.

41. Spring Session

Spring Boot provides Spring Session auto-configuration for a wide range of stores:

- JDBC
- Redis
- Hazelcast
- HashMap

If Spring Session is available, you must choose the [StoreType](#) that you wish to use to store the sessions. For instance to use JDBC as backend store, you'd configure your application as follows:

```
spring.session.store-type=jdbc
```

Tip

You can disable Spring Session by setting the `store-type` to `none`.

Each store has specific additional settings. For instance it is possible to customize the name of the table for the jdbc store:

```
spring.session.jdbc.table-name=SESSIONS
```

42. Monitoring and management over JMX

Java Management Extensions (JMX) provide a standard mechanism to monitor and manage applications. By default Spring Boot will create an `MBeanServer` with bean id 'mbeanServer' and expose any of your beans that are annotated with Spring JMX annotations (`@ManagedResource`, `@ManagedAttribute`, `@ManagedOperation`).

See the [JmxAutoConfiguration](#) class for more details.

43. Testing

Spring Boot provides a number of utilities and annotations to help when testing your application. Test support is provided by two modules; `spring-boot-test` contains core items, and `spring-boot-test-autoconfigure` supports auto-configuration for tests.

Most developers will just use the `spring-boot-starter-test` ‘Starter’ which imports both Spring Boot test modules as well as JUnit, AssertJ, Hamcrest and a number of other useful libraries.

43.1 Test scope dependencies

If you use the `spring-boot-starter-test` ‘Starter’ (in the `test` scope), you will find the following provided libraries:

- [JUnit](#) — The de-facto standard for unit testing Java applications.
- [Spring Test](#) & Spring Boot Test — Utilities and integration test support for Spring Boot applications.
- [AssertJ](#) — A fluent assertion library.
- [Hamcrest](#) — A library of matcher objects (also known as constraints or predicates).
- [Mockito](#) — A Java mocking framework.
- [JSONassert](#) — An assertion library for JSON.
- [JsonPath](#) — XPath for JSON.

These are common libraries that we generally find useful when writing tests. You are free to add additional test dependencies of your own if these don’t suit your needs.

43.2 Testing Spring applications

One of the major advantages of dependency injection is that it should make your code easier to unit test. You can simply instantiate objects using the `new` operator without even involving Spring. You can also use *mock objects* instead of real dependencies.

Often you need to move beyond ‘unit testing’ and start ‘integration testing’ (with a Spring `ApplicationContext` actually involved in the process). It’s useful to be able to perform integration testing without requiring deployment of your application or needing to connect to other infrastructure.

The Spring Framework includes a dedicated test module for just such integration testing. You can declare a dependency directly to `org.springframework:spring-test` or use the `spring-boot-starter-test` ‘Starter’ to pull it in transitively.

If you have not used the `spring-test` module before you should start by reading the [relevant section](#) of the Spring Framework reference documentation.

43.3 Testing Spring Boot applications

A Spring Boot application is just a Spring `ApplicationContext`, so nothing very special has to be done to test it beyond what you would normally do with a vanilla Spring context. One thing to watch out for though is that the external properties, logging and other features of Spring Boot are only installed in the context by default if you use `SpringApplication` to create it.

Spring Boot provides a `@SpringBootTest` annotation which can be used as an alternative to the standard `spring-test` `@ContextConfiguration` annotation when you need Spring Boot features. The annotation works by creating the `ApplicationContext` used in your tests via `SpringApplication`.

You can use the `webEnvironment` attribute of `@SpringBootTest` to further refine how your tests will run:

- **MOCK** — Loads a `WebApplicationContext` and provides a mock servlet environment. Embedded servlet containers are not started when using this annotation. If servlet APIs are not on your classpath this mode will transparently fallback to creating a regular non-web `ApplicationContext`. Can be used in conjunction with `@AutoConfigureMockMvc` for `MockMvc`-based testing of your application.
- **RANDOM_PORT** — Loads an `ServletWebServerApplicationContext` and provides a real servlet environment. Embedded servlet containers are started and listening on a random port.
- **DEFINED_PORT** — Loads an `ServletWebServerApplicationContext` and provides a real servlet environment. Embedded servlet containers are started and listening on a defined port (i.e from your `application.properties` or on the default port 8080).
- **NONE** — Loads an `ApplicationContext` using `SpringApplication` but does not provide *any* servlet environment (mock or otherwise).

Note

If your test is `@Transactional`, it will rollback the transaction at the end of each test method by default. If you're using this arrangement in combination with either `RANDOM_PORT` or `DEFINED_PORT`, any transaction initiated on the server won't rollback as the test is running in a different thread than the server processing.

Note

In addition to `@SpringBootTest` a number of other annotations are also provided for testing more specific slices of an application. See below for details.

Tip

Don't forget to also add `@RunWith(SpringRunner.class)` to your test, otherwise the annotations will be ignored.

Detecting test configuration

If you're familiar with the Spring Test Framework, you may be used to using `@ContextConfiguration(classes=...)` in order to specify which Spring `@Configuration` to load. Alternatively, you might have often used nested `@Configuration` classes within your test.

When testing Spring Boot applications this is often not required. Spring Boot's `@*Test` annotations will search for your primary configuration automatically whenever you don't explicitly define one.

The search algorithm works up from the package that contains the test until it finds a `@SpringBootApplication` or `@SpringBootConfiguration` annotated class. As long as you've [structured your code](#) in a sensible way your main configuration is usually found.

If you want to customize the primary configuration, you can use a nested `@TestConfiguration` class. Unlike a nested `@Configuration` class which would be used instead of your application's primary configuration, a nested `@TestConfiguration` class will be used in addition to your application's primary configuration.

Note

Spring's test framework will cache application contexts between tests. Therefore, as long as your tests share the same configuration (no matter how it's discovered), the potentially time consuming process of loading the context will only happen once.

Excluding test configuration

If your application uses component scanning, for example if you use `@SpringBootApplication` or `@ComponentScan`, you may find top-level configuration classes created only for specific tests accidentally get picked up everywhere.

As we [have seen above](#), `@TestConfiguration` can be used on an inner class of a test to customize the primary configuration. When placed on a top-level class, `@TestConfiguration` indicates that classes in `src/test/java` should not be picked up by scanning. You can then import that class explicitly where it is required:

```
@RunWith(SpringRunner.class)
@SpringBootTest
@Import(MyTestsConfiguration.class)
public class MyTests {

    @Test
    public void exampleTest() {
        ...
    }
}
```

Note

If you directly use `@ComponentScan` (i.e. not via `@SpringBootApplication`) you will need to register the `TypeExcludeFilter` with it. See [the Javadoc](#) for details.

Working with random ports

If you need to start a full running server for tests, we recommend that you use random ports. If you use `@SpringBootTest(webEnvironment=WebEnvironment.RANDOM_PORT)` an available port will be picked at random each time your test runs.

The `@LocalServerPort` annotation can be used to [inject the actual port used](#) into your test. For convenience, tests that need to make REST calls to the started server can additionally `@Autowire` a `TestRestTemplate` which will resolve relative links to the running server.

```
import org.junit.Test;
import org.junit.runner.RunWith;

import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.context.SpringBootTest;
import org.springframework.boot.test.context.SpringBootTest.WebEnvironment;
import org.springframework.boot.test.web.client.TestRestTemplate;
import org.springframework.test.context.junit4.SpringRunner;
```

```

import static org.assertj.core.api.Assertions.assertThat;

@RunWith(SpringRunner.class)
@SpringBootTest(webEnvironment = WebEnvironment.RANDOM_PORT)
public class RandomPortExampleTests {

    @Autowired
    private TestRestTemplate restTemplate;

    @Test
    public void exampleTest() {
        String body = this.restTemplate.getForObject("/", String.class);
        assertThat(body).isEqualTo("Hello World");
    }
}

```

Mocking and spying beans

It's sometimes necessary to mock certain components within your application context when running tests. For example, you may have a facade over some remote service that's unavailable during development. Mocking can also be useful when you want to simulate failures that might be hard to trigger in a real environment.

Spring Boot includes a `@MockBean` annotation that can be used to define a Mockito mock for a bean inside your `ApplicationContext`. You can use the annotation to add new beans, or replace a single existing bean definition. The annotation can be used directly on test classes, on fields within your test, or on `@Configuration` classes and fields. When used on a field, the instance of the created mock will also be injected. Mock beans are automatically reset after each test method.

Note

This feature is automatically enabled as long as your test uses one of Spring Boot's test annotations (i.e. `@SpringBootTest`). To use this feature with a different arrangement, a listener will need to be added explicitly:

```
@TestExecutionListeners(MockitoTestExecutionListener.class)
```

Here's a typical example where we replace an existing `RemoteService` bean with a mock implementation:

```

import org.junit.*;
import org.junit.runner.*;
import org.springframework.beans.factory.annotation.*;
import org.springframework.boot.test.context.*;
import org.springframework.boot.test.mock.mockito.*;
import org.springframework.test.context.junit4.*;

import static org.assertj.core.api.Assertions.*;
import static org.mockito.BDDMockito.*;

@RunWith(SpringRunner.class)
@SpringBootTest
public class MyTests {

    @MockBean
    private RemoteService remoteService;

    @Autowired
    private Reverser reverser;
}

```

```

@Test
public void exampleTest() {
    // RemoteService has been injected into the reverser bean
    given(this.remoteService.someCall()).willReturn("mock");
    String reverse = reverser.reverseSomeCall();
    assertThat(reverse).isEqualTo("kcom");
}
}

```

Additionally you can also use `@SpyBean` to wrap any existing bean with a Mockito `spy`. See the Javadoc for full details.

Auto-configured tests

Spring Boot's auto-configuration system works well for applications, but can sometimes be a little too much for tests. It's often helpful to load only the parts of the configuration that are required to test a 'slice' of your application. For example, you might want to test that Spring MVC controllers are mapping URLs correctly, and you don't want to involve database calls in those tests; or you *might be wanting* to test JPA entities, and you're not interested in web layer when those tests run.

The `spring-boot-test-autoconfigure` module includes a number of annotations that can be used to automatically configure such 'slices'. Each of them works in a similar way, providing a `@...Test` annotation that loads the `ApplicationContext` and one or more `@AutoConfigure...` annotations that can be used to customize auto-configuration settings.

Note

Each slice loads a very restricted set of auto-configuration classes. If you need to exclude one of them, most `@...Test` annotations provide an `excludeAutoConfiguration` attribute. Alternatively, you can use `@ImportAutoConfiguration#exclude`.

Tip

It's also possible to use the `@AutoConfigure...` annotations with the standard `@SpringBootTest` annotation. You can use this combination if you're not interested in 'slicing' your application but you want some of the auto-configured test beans.

Auto-configured JSON tests

To test that Object JSON serialization and deserialization is working as expected you can use the `@JsonTest` annotation. `@JsonTest` will auto-configure Jackson `ObjectMapper`, any `@JsonComponent` beans and any Jackson `Modules`. It also configures `Gson` if you happen to be using that instead of, or as well as, Jackson. If you need to configure elements of the auto-configuration you can use the `@AutoConfigureJsonTesters` annotation.

Spring Boot includes AssertJ based helpers that work with the `JSONassert` and `JsonPath` libraries to check that JSON is as expected. The `JacksonTester`, `GsonTester` and `BasicJsonTester` classes can be used for Jackson, Gson and Strings respectively. Any helper fields on the test class can be `@Autowired` when using `@JsonTest`.

```

import org.junit.*;
import org.junit.runner.*;
import org.springframework.beans.factory.annotation.*;
import org.springframework.boot.test.autoconfigure.json.*;

```

```

import org.springframework.boot.test.context.*;
import org.springframework.boot.test.json.*;
import org.springframework.test.context.junit4.*;

import static org.assertj.core.api.Assertions.*;

@RunWith(SpringRunner.class)
@JacksonTest
public class MyJsonTests {

    @Autowired
    private JacksonTester<VehicleDetails> json;

    @Test
    public void testSerialize() throws Exception {
        VehicleDetails details = new VehicleDetails("Honda", "Civic");
        // Assert against a `.json` file in the same package as the test
        assertThat(this.json.write(details)).isEqualToJson("expected.json");
        // Or use JSON path based assertions
        assertThat(this.json.write(details)).hasJsonPathStringValue("@.make");
        assertThat(this.json.write(details).extractingJsonPathStringValue("@.make")
            .isEqualTo("Honda"));
    }

    @Test
    public void testDeserialize() throws Exception {
        String content = "{\"make\":\"Ford\", \"model\":\"Focus\"}";
        assertThat(this.json.parse(content))
            .isEqualTo(new VehicleDetails("Ford", "Focus"));
        assertThat(this.json.parseObject(content).getMake()).isEqualTo("Ford");
    }
}

```

Note

JSON helper classes can also be used directly in standard unit tests. Simply call the `initFields` method of the helper in your `@Before` method if you aren't using `@JacksonTest`.

A list of the auto-configuration that is enabled by `@JacksonTest` can be [found in the appendix](#).

Auto-configured Spring MVC tests

To test Spring MVC controllers are working as expected you can use the `@WebMvcTest` annotation. `@WebMvcTest` will auto-configure the Spring MVC infrastructure and limit scanned beans to `@Controller`, `@ControllerAdvice`, `@JsonComponent`, `Filter`, `WebMvcConfigurer` and `HandlerMethodArgumentResolver`. Regular `@Component` beans will not be scanned when using this annotation.

Often `@WebMvcTest` will be limited to a single controller and used in combination with `@MockBean` to provide mock implementations for required collaborators.

`@WebMvcTest` also auto-configures `MockMvc`. `MockMvc` offers a powerful way to quickly test MVC controllers without needing to start a full HTTP server.

Tip

You can also auto-configure `MockMvc` in a non-`@WebMvcTest` (e.g. `SpringBootTest`) by annotating it with `@AutoConfigureMockMvc`.

```
import org.junit.*;
```

```

import org.junit.runner.*;
import org.springframework.beans.factory.annotation.*;
import org.springframework.boot.test.autoconfigure.web.servlet.*;
import org.springframework.boot.test.mock.mockito.*;

import static org.assertj.core.api.Assertions.*;
import static org.mockito.BDDMockito.*;
import static org.springframework.test.web.servlet.request.MockMvcRequestBuilders.*;
import static org.springframework.test.web.servlet.result.MockMvcResultMatchers.*;

@RunWith(SpringRunner.class)
@WebMvcTest(UserVehicleController.class)
public class MyControllerTests {

    @Autowired
    private MockMvc mvc;

    @MockBean
    private UserVehicleService userVehicleService;

    @Test
    public void testExample() throws Exception {
        given(this.userVehicleService.getVehicleDetails("sboot"))
            .willReturn(new VehicleDetails("Honda", "Civic"));
        this.mvc.perform(get("/sboot/vehicle").accept(MediaType.TEXT_PLAIN))
            .andExpect(status().isOk()).andExpect(content().string("Honda Civic"));
    }
}

```

Tip

If you need to configure elements of the auto-configuration (for example when servlet filters should be applied) you can use attributes in the `@AutoConfigureMockMvc` annotation.

If you use `HtmlUnit` or `Selenium`, auto-configuration will also provide an `HTMLUnit WebClient` bean and/or a `WebDriver` bean. Here is an example that uses `HtmlUnit`:

```

import com.gargoylesoftware.htmlunit.*;
import org.junit.*;
import org.junit.runner.*;
import org.springframework.beans.factory.annotation.*;
import org.springframework.boot.test.autoconfigure.web.servlet.*;
import org.springframework.boot.test.mock.mockito.*;

import static org.assertj.core.api.Assertions.*;
import static org.mockito.BDDMockito.*;

@RunWith(SpringRunner.class)
@WebMvcTest(UserVehicleController.class)
public class MyHtmlUnitTests {

    @Autowired
    private WebClient webClient;

    @MockBean
    private UserVehicleService userVehicleService;

    @Test
    public void testExample() throws Exception {
        given(this.userVehicleService.getVehicleDetails("sboot"))
            .willReturn(new VehicleDetails("Honda", "Civic"));
        HtmlPage page = this.webClient.getPage("/sboot/vehicle.html");
        assertThat(page.getBody().getTextContent()).isEqualTo("Honda Civic");
    }
}

```

Note

By default Spring Boot will put `WebDriver` beans in a special “scope” to ensure that the driver is quit after each test, and that a new instance is injected. If you don't want this behavior you can add `@Scope("singleton")` to your `WebDriver` `@Bean` definition.

A list of the auto-configuration that is enabled by `@WebMvcTest` can be [found in the appendix](#).

Auto-configured Spring WebFlux tests

To test Spring WebFlux controllers are working as expected you can use the `@WebFluxTest` annotation. `@WebFluxTest` will auto-configure the Spring WebFlux infrastructure and limit scanned beans to `@Controller`, `@ControllerAdvice`, `@JsonComponent`, and `WebFluxConfigurer`. Regular `@Component` beans will not be scanned when using this annotation.

Often `@WebFluxTest` will be limited to a single controller and used in combination with `@MockBean` to provide mock implementations for required collaborators.

`@WebFluxTest` also auto-configures `WebTestClient`, which offers a powerful way to quickly test WebFlux controllers without needing to start a full HTTP server.

Tip

You can also auto-configure `WebTestClient` in a non-`@WebFluxTest` (e.g. `SpringBootTest`) by annotating it with `@AutoConfigureWebTestClient`.

```
import org.junit.Test;
import org.junit.runner.RunWith;

import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.web.reactive.WebFluxTest;
import org.springframework.http.MediaType;
import org.springframework.test.context.junit4.SpringRunner;
import org.springframework.test.web.reactive.server.WebTestClient;

@RunWith(SpringRunner.class)
@WebFluxTest(UserVehicleController.class)
public class MyControllerTests {

    @Autowired
    private WebTestClient webClient;

    @MockBean
    private UserVehicleService userVehicleService;

    @Test
    public void testExample() throws Exception {
        given(this.userVehicleService.getVehicleDetails("sboot"))
            .willReturn(new VehicleDetails("Honda", "Civic"));
        this.webClient.get().uri("/sboot/vehicle").accept(MediaType.TEXT_PLAIN)
            .exchange()
            .expectStatus().isOk()
            .expectBody(String.class).isEqualTo("Honda Civic");
    }
}
```

A list of the auto-configuration that is enabled by `@WebFluxTest` can be [found in the appendix](#).

Auto-configured Data JPA tests

`@DataJpaTest` can be used if you want to test JPA applications. By default it will configure an in-memory embedded database, scan for `@Entity` classes and configure Spring Data JPA repositories. Regular `@Component` beans will not be loaded into the `ApplicationContext`.

Data JPA tests are transactional and rollback at the end of each test by default, see the [relevant section](#) in the Spring Reference Documentation for more details. If that's not what you want, you can disable transaction management for a test or for the whole class as follows:

```
import org.junit.Test;
import org.junit.runner.RunWith;
import org.springframework.boot.test.autoconfigure.orm.jpa.DataJpaTest;
import org.springframework.test.context.junit4.SpringRunner;
import org.springframework.transaction.annotation.Propagation;
import org.springframework.transaction.annotation.Transactional;

@RunWith(SpringRunner.class)
@DataJpaTest
@Transactional(propagation = Propagation.NOT_SUPPORTED)
public class ExampleNonTransactionalTests {

}
```

Data JPA tests may also inject a [TestEntityManager](#) bean which provides an alternative to the standard JPA `EntityManager` specifically designed for tests. If you want to use `TestEntityManager` outside of `@DataJpaTests` you can also use the `@AutoConfigureTestEntityManager` annotation. A `JdbcTemplate` is also available if you need that.

```
import org.junit.*;
import org.junit.runner.*;
import org.springframework.boot.test.autoconfigure.orm.jpa.*;

import static org.assertj.core.api.Assertions.*;

@RunWith(SpringRunner.class)
@DataJpaTest
public class ExampleRepositoryTests {

    @Autowired
    private TestEntityManager entityManager;

    @Autowired
    private UserRepository repository;

    @Test
    public void testExample() throws Exception {
        this.entityManager.persist(new User("sboot", "1234"));
        User user = this.repository.findByUsername("sboot");
        assertThat(user.getUsername()).isEqualTo("sboot");
        assertThat(user.getVin()).isEqualTo("1234");
    }

}
```

In-memory embedded databases generally work well for tests since they are fast and don't require any developer installation. If, however, you prefer to run tests against a real database you can use the `@AutoConfigureTestDatabase` annotation:

```
@RunWith(SpringRunner.class)
@DataJpaTest
@AutoConfigureTestDatabase(replace=Replace.NONE)
public class ExampleRepositoryTests {

}
```

```
// ...
}
```

A list of the auto-configuration that is enabled by `@DataJpaTest` can be [found in the appendix](#).

Auto-configured JDBC tests

`@JdbcTest` is similar to `@DataJpaTest` but for pure jdbc-related tests. By default it will also configure an in-memory embedded database and a `JdbcTemplate`. Regular `@Component` beans will not be loaded into the `ApplicationContext`.

JDBC tests are transactional and rollback at the end of each test by default, see the [relevant section](#) in the Spring Reference Documentation for more details. If that's not what you want, you can disable transaction management for a test or for the whole class as follows:

```
import org.junit.Test;
import org.junit.runner.RunWith;
import org.springframework.boot.test.autoconfigure.jdbc.JdbcTest;
import org.springframework.test.context.junit4.SpringRunner;
import org.springframework.transaction.annotation.Propagation;
import org.springframework.transaction.annotation.Transactional;

@RunWith(SpringRunner.class)
@JdbcTest
@Transactional(propagation = Propagation.NOT_SUPPORTED)
public class ExampleNonTransactionalTests {

}
```

If you prefer your test to run against a real database, you can use the `@AutoConfigureTestDatabase` annotation the same way as for `DataJpaTest`.

A list of the auto-configuration that is enabled by `@JdbcTest` can be [found in the appendix](#).

Auto-configured jOOQ tests

`@JooqTest` can be used in a similar fashion as `@JdbcTest` but for jOOQ related tests. As jOOQ relies heavily on a Java-based schema that corresponds with the database schema, the existing `DataSource` will be used. If you want to replace it by an in-memory database you can use `@AutoConfigureTestDatabase` to override those settings.

`@JooqTest` will configure a `DSLContext`. Regular `@Component` beans will not be loaded into the `ApplicationContext`:

```
import org.jooq.DSLContext;
import org.junit.Test;
import org.junit.runner.RunWith;
import org.springframework.boot.test.autoconfigure.jooq.JooqTest;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@JooqTest
public class ExampleJooqTests {

    @Autowired
    private DSLContext dslContext;

}
```

JOOQ tests are transactional and rollback at the end of each test by default. If that's not what you want, you can disable transaction management for a test or for the whole test class as [shown in the example above](#).

A list of the auto-configuration that is enabled by `@JooqTest` can be [found in the appendix](#).

Auto-configured Data MongoDB tests

`@DataMongoTest` can be used if you want to test MongoDB applications. By default, it will configure an in-memory embedded MongoDB (if available), configure a `MongoTemplate`, scan for `@Document` classes and configure Spring Data MongoDB repositories. Regular `@Component` beans will not be loaded into the `ApplicationContext`:

```
import org.junit.runner.RunWith;
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.data.mongo.DataMongoTest;
import org.springframework.data.mongodb.core.MongoTemplate;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataMongoTest
public class ExampleDataMongoTests {

    @Autowired
    private MongoTemplate mongoTemplate;

    //
}
```

In-memory embedded MongoDB generally works well for tests since it is fast and doesn't require any developer installation. If, however, you prefer to run tests against a real MongoDB server you should exclude the embedded MongoDB auto-configuration:

```
import org.junit.runner.RunWith;
import org.springframework.boot.autoconfigure.mongo.embedded.EmbeddedMongoAutoConfiguration;
import org.springframework.boot.test.autoconfigure.data.mongo.DataMongoTest;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataMongoTest(excludeAutoConfiguration = EmbeddedMongoAutoConfiguration.class)
public class ExampleDataMongoNonEmbeddedTests {

}
```

A list of the auto-configuration that is enabled by `@DataMongoTest` can be [found in the appendix](#).

Auto-configured Data Neo4j tests

`@DataNeo4jTest` can be used if you want to test Neo4j applications. By default, it will use an in-memory embedded Neo4j (if the embedded driver is available), scan for `@NodeEntity` classes and configure Spring Data Neo4j repositories. Regular `@Component` beans will not be loaded into the `ApplicationContext`:

```
import org.junit.runner.RunWith;
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.data.neo4j.DataNeo4jTest;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataNeo4jTest
public class ExampleDataNeo4jTests {

    @Autowired
    private YourRepository repository;

    //
}
```

```
}

```

Data Neo4j tests are transactional and rollback at the end of each test by default, see the [relevant section](#) in the Spring Reference Documentation for more details. If that's not what you want, you can disable transaction management for a test or for the whole class as follows:

```
import org.junit.Test;
import org.junit.runner.RunWith;
import org.springframework.boot.test.autoconfigure.data.neo4j.DataNeo4jTest;
import org.springframework.test.context.junit4.SpringRunner;
import org.springframework.transaction.annotation.Propagation;
import org.springframework.transaction.annotation.Transactional;

@RunWith(SpringRunner.class)
@DataNeo4jTest
@Transactional(propagation = Propagation.NOT_SUPPORTED)
public class ExampleNonTransactionalTests {

}

```

A list of the auto-configuration that is enabled by `@DataNeo4jTest` can be [found in the appendix](#).

Auto-configured Data Redis tests

`@DataRedisTest` can be used if you want to test Redis applications. By default, it will scan for `@RedisHash` classes and configure Spring Data Redis repositories. Regular `@Component` beans will not be loaded into the `ApplicationContext`:

```
import org.junit.runner.RunWith;
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.data.redis.DataRedisTest;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataRedisTest
public class ExampleDataRedisTests {

    @Autowired
    private YourRepository repository;

    //
}

```

A list of the auto-configuration that is enabled by `@DataRedisTest` can be [found in the appendix](#).

Auto-configured Data LDAP tests

`@DataLdapTest` can be used if you want to test LDAP applications. By default, it will configure an in-memory embedded LDAP (if available), a `LdapTemplate`, scan for `@Entry` classes and configure Spring Data LDAP repositories. Regular `@Component` beans will not be loaded into the `ApplicationContext`:

```
import org.junit.runner.RunWith;
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.data.ldap.DataLdapTest;
import org.springframework.ldap.core.LdapTemplate;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataLdapTest
public class ExampleDataLdapTests {

    @Autowired

```

```
private LdapTemplate ldapTemplate;

//
}
```

In-memory embedded LDAP generally works well for tests since it is fast and doesn't require any developer installation. If, however, you prefer to run tests against a real LDAP server you should exclude the embedded LDAP auto-configuration:

```
import org.junit.runner.RunWith;
import org.springframework.boot.autoconfigure.ldap.embedded.EmbeddedLdapAutoConfiguration;
import org.springframework.boot.test.autoconfigure.data.ldap.DataLdapTest;
import org.springframework.test.context.junit4.SpringRunner;

@RunWith(SpringRunner.class)
@DataLdapTest(excludeAutoConfiguration = EmbeddedLdapAutoConfiguration.class)
public class ExampleDataLdapNonEmbeddedTests {

}
```

A list of the auto-configuration that is enabled by `@DataLdapTest` can be [found in the appendix](#).

Auto-configured REST clients

The `@RestClientTest` annotation can be used if you want to test REST clients. By default it will auto-configure Jackson and GSON support, configure a `RestTemplateBuilder` and add support for `MockRestServiceServer`. The specific beans that you want to test should be specified using `value` or `components` attribute of `@RestClientTest`:

```
@RunWith(SpringRunner.class)
@RestClientTest(RemoteVehicleDetailsService.class)
public class ExampleRestClientTest {

    @Autowired
    private RemoteVehicleDetailsService service;

    @Autowired
    private MockRestServiceServer server;

    @Test
    public void getVehicleDetailsWhenResultIsSuccessShouldReturnDetails()
        throws Exception {
        this.server.expect(requestTo("/greet/details"))
            .andRespond(withSuccess("hello", MediaType.TEXT_PLAIN));
        String greeting = this.service.callRestService();
        assertThat(greeting).isEqualTo("hello");
    }

}
```

A list of the auto-configuration that is enabled by `@RestClientTest` can be [found in the appendix](#).

Auto-configured Spring REST Docs tests

The `@AutoConfigureRestDocs` annotation can be used if you want to use Spring REST Docs in your tests. It will automatically configure `MockMvc` to use Spring REST Docs and remove the need for Spring REST Docs' JUnit rule.

```
import org.junit.Test;
import org.junit.runner.RunWith;

import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.test.autoconfigure.web.servlet.WebMvcTest;
```

```

import org.springframework.http.MediaType;
import org.springframework.test.context.junit4.SpringRunner;
import org.springframework.test.web.servlet.MockMvc;

import static org.springframework.restdocs.mockmvc.MockMvcRestDocumentation.document;
import static org.springframework.test.web.servlet.request.MockMvcRequestBuilders.get;
import static org.springframework.test.web.servlet.result.MockMvcResultMatchers.*;

@RunWith(SpringRunner.class)
@WebMvcTest(UserController.class)
@AutoConfigureRestDocs
public class UserDocumentationTests {

    @Autowired
    private MockMvc mvc;

    @Test
    public void listUsers() throws Exception {
        this.mvc.perform(get("/users").accept(MediaType.TEXT_PLAIN))
            .andExpect(status().isOk())
            .andDo(document("list-users"));
    }
}

```

`@AutoConfigureRestDocs` can be used to override the default output directory (target/generated-snippets if you are using Maven or build/generated-snippets if you are using Gradle). It can also be used to configure the host, scheme, and port that will appear in any documented URIs. If you require more control over Spring REST Docs' configuration a `RestDocsMockMvcConfigurationCustomizer` bean can be used:

```

@TestConfiguration
static class CustomizationConfiguration
    implements RestDocsMockMvcConfigurationCustomizer {

    @Override
    public void customize(MockMvcRestDocumentationConfigurer configurer) {
        configurer.snippets().withTemplateFormat(TemplateFormats.markdown());
    }
}

```

If you want to make use of Spring REST Docs' support for a parameterized output directory, you can create a `RestDocumentationResultHandler` bean. The auto-configuration will call `alwaysDo` with this result handler, thereby causing each `MockMvc` call to automatically generate the default snippets:

```

@TestConfiguration
static class ResultHandlerConfiguration {

    @Bean
    public RestDocumentationResultHandler restDocumentation() {
        return MockMvcRestDocumentation.document("{method-name}");
    }
}

```

Using Spock to test Spring Boot applications

If you wish to use Spock to test a Spring Boot application you should add a dependency on Spock's `spock-spring` module to your application's build. `spock-spring` integrates Spring's test framework into Spock. It is recommended that you use Spock 1.1 or later to benefit from a number of recent improvements to Spock's Spring Framework and Spring Boot integration. Please refer to [the documentation for Spock's Spring module](#) for further details.

43.4 Test utilities

A few test utility classes are packaged as part of `spring-boot` that are generally useful when testing your application.

ConfigFileApplicationContextInitializer

`ConfigFileApplicationContextInitializer` is an `ApplicationContextInitializer` that can apply to your tests to load Spring Boot `application.properties` files. You can use this when you don't need the full features provided by `@SpringBootTest`.

```
@ContextConfiguration(classes = Config.class,
    initializers = ConfigFileApplicationContextInitializer.class)
```

Note

Using `ConfigFileApplicationContextInitializer` alone won't provide support for `@Value("${...}")` injection. Its only job is to ensure that `application.properties` files are loaded into Spring's Environment. For `@Value` support you need to either additionally configure a `PropertySourcesPlaceholderConfigurer` or use `@SpringBootTest` where one will be auto-configured for you.

EnvironmentTestUtils

`EnvironmentTestUtils` allows you to quickly add properties to a `ConfigurableEnvironment` or `ConfigurableApplicationContext`. Simply call it with `key=value` strings:

```
EnvironmentTestUtils.addEnvironment(env, "org=Spring", "name=Boot");
```

OutputCapture

`OutputCapture` is a JUnit Rule that you can use to capture `System.out` and `System.err` output. Simply declare the capture as a `@Rule` then use `toString()` for assertions:

```
import org.junit.Rule;
import org.junit.Test;
import org.springframework.boot.test.rule.OutputCapture;

import static org.hamcrest.Matchers.*;
import static org.junit.Assert.*;

public class MyTest {

    @Rule
    public OutputCapture capture = new OutputCapture();

    @Test
    public void testName() throws Exception {
        System.out.println("Hello World!");
        assertThat(capture.toString(), containsString("World"));
    }

}
```

TestRestTemplate

`TestRestTemplate` is a convenience alternative to Spring's `RestTemplate` that is useful in integration tests. You can get a vanilla template or one that sends Basic HTTP authentication (with a

username and password). In either case the template will behave in a test-friendly way by not throwing exceptions on server-side errors. It is recommended, but not mandatory, to use Apache HTTP Client (version 4.3.2 or better), and if you have that on your classpath the `TestRestTemplate` will respond by configuring the client appropriately. If you do use Apache's HTTP client some additional test-friendly features will be enabled:

- Redirects will not be followed (so you can assert the response location)
- Cookies will be ignored (so the template is stateless)

`TestRestTemplate` can be instantiated directly in your integration tests:

```
public class MyTest {

    private TestRestTemplate template = new TestRestTemplate();

    @Test
    public void testRequest() throws Exception {
        HttpHeaders headers = template.getForEntity("http://myhost.com/example",
String.class).getHeaders();
        assertThat(headers.getLocation().toString(), containsString("myotherhost"));
    }

}
```

Alternatively, if you are using the `@SpringBootTest` annotation with `WebEnvironment.RANDOM_PORT` or `WebEnvironment.DEFINED_PORT`, you can just inject a fully configured `TestRestTemplate` and start using it. If necessary, additional customizations can be applied via the `RestTemplateBuilder` bean. Any URLs that do not specify a host and port will automatically connect to the embedded server:

```
@RunWith(SpringRunner.class)
@SpringBootTest
public class MyTest {

    @Autowired
    private TestRestTemplate template;

    @Test
    public void testRequest() throws Exception {
        HttpHeaders headers = template.getForEntity("/example", String.class).getHeaders();
        assertThat(headers.getLocation().toString(), containsString("myotherhost"));
    }

    @TestConfiguration
    static class Config {

        @Bean
        public RestTemplateBuilder restTemplateBuilder() {
            return new RestTemplateBuilder()
                .additionalMessageConverters(...)
                .customizers(...);
        }

    }

}
```

44. WebSockets

Spring Boot provides WebSockets auto-configuration for embedded Tomcat (8 and 7), Jetty 9 and Undertow. If you're deploying a war file to a standalone container, Spring Boot assumes that the container will be responsible for the configuration of its WebSocket support.

Spring Framework provides [rich WebSocket support](#) that can be easily accessed via the `spring-boot-starter-websocket` module.

45. Web Services

Spring Boot provides Web Services auto-configuration so that all is required is defining your Endpoints.

The [Spring Web Services features](#) can be easily accessed via the `spring-boot-starter-webservices` module.

46. Creating your own auto-configuration

If you work in a company that develops shared libraries, or if you work on an open-source or commercial library, you might want to develop your own auto-configuration. Auto-configuration classes can be bundled in external jars and still be picked-up by Spring Boot.

Auto-configuration can be associated to a "starter" that provides the auto-configuration code as well as the typical libraries that you would use with it. We will first cover what you need to know to build your own auto-configuration and we will move on to the [typical steps required to create a custom starter](#).

Tip

A [demo project](#) is available to showcase how you can create a starter step by step.

46.1 Understanding auto-configured beans

Under the hood, auto-configuration is implemented with standard `@Configuration` classes. Additional `@Conditional` annotations are used to constrain when the auto-configuration should apply. Usually auto-configuration classes use `@ConditionalOnClass` and `@ConditionalOnMissingBean` annotations. This ensures that auto-configuration only applies when relevant classes are found and when you have not declared your own `@Configuration`.

You can browse the source code of [spring-boot-autoconfigure](#) to see the `@Configuration` classes that we provide (see the [META-INF/spring.factories](#) file).

46.2 Locating auto-configuration candidates

Spring Boot checks for the presence of a `META-INF/spring.factories` file within your published jar. The file should list your configuration classes under the `EnableAutoConfiguration` key.

```
org.springframework.boot.autoconfigure.EnableAutoConfiguration=\
com.mycorp.libx.autoconfigure.LibXAutoConfiguration,\
com.mycorp.libx.autoconfigure.LibXWebAutoConfiguration
```

You can use the `@AutoConfigureAfter` or `@AutoConfigureBefore` annotations if your configuration needs to be applied in a specific order. For example, if you provide web-specific configuration, your class may need to be applied after `WebMvcAutoConfiguration`.

If you want to order certain auto-configurations that shouldn't have any direct knowledge of each other, you can also use `@AutoConfigureOrder`. That annotation has the same semantic as the regular `@Order` annotation but provides a dedicated order for auto-configuration classes.

Note

Auto-configurations have to be loaded that way *only*. Make sure that they are defined in a specific package space and that they are never the target of component scan in particular.

46.3 Condition annotations

You almost always want to include one or more `@Conditional` annotations on your auto-configuration class. The `@ConditionalOnMissingBean` is one common example that is used to allow developers to 'override' auto-configuration if they are not happy with your defaults.

Spring Boot includes a number of `@Conditional` annotations that you can reuse in your own code by annotating `@Configuration` classes or individual `@Bean` methods.

Class conditions

The `@ConditionalOnClass` and `@ConditionalOnMissingClass` annotations allows configuration to be included based on the presence or absence of specific classes. Due to the fact that annotation metadata is parsed using [ASM](#) you can actually use the `value` attribute to refer to the real class, even though that class might not actually appear on the running application classpath. You can also use the `name` attribute if you prefer to specify the class name using a `String` value.

Tip

If you are using `@ConditionalOnClass` or `@ConditionalOnMissingClass` as a part of a meta-annotation to compose your own composed annotations you must use `name` as referring to the class in such a case is not handled.

Bean conditions

The `@ConditionalOnBean` and `@ConditionalOnMissingBean` annotations allow a bean to be included based on the presence or absence of specific beans. You can use the `value` attribute to specify beans by type, or `name` to specify beans by name. The `search` attribute allows you to limit the `ApplicationContext` hierarchy that should be considered when searching for beans.

When placed on a `@Bean` method, the target type defaults to the return type of the method, for instance:

```
@Configuration
public class MyAutoConfiguration {

    @Bean
    @ConditionalOnMissingBean
    public MyService myService() { ... }

}
```

In the example above, the `myService` bean is going to be created if no bean of type `MyService` is already contained in the `ApplicationContext`.

Tip

You need to be very careful about the order that bean definitions are added as these conditions are evaluated based on what has been processed so far. For this reason, we recommend only using `@ConditionalOnBean` and `@ConditionalOnMissingBean` annotations on auto-configuration classes (since these are guaranteed to load after any user-defined beans definitions have been added).

Note

`@ConditionalOnBean` and `@ConditionalOnMissingBean` do not prevent `@Configuration` classes from being created. Using these conditions at the class level is equivalent to marking each contained `@Bean` method with the annotation.

Property conditions

The `@ConditionalOnProperty` annotation allows configuration to be included based on a Spring Environment property. Use the `prefix` and `name` attributes to specify the property that should be checked. By default any property that exists and is not equal to `false` will be matched. You can also create more advanced checks using the `havingValue` and `matchIfMissing` attributes.

Resource conditions

The `@ConditionalOnResource` annotation allows configuration to be included only when a specific resource is present. Resources can be specified using the usual Spring conventions, for example, `file:/home/user/test.dat`.

Web application conditions

The `@ConditionalOnWebApplication` and `@ConditionalOnNotWebApplication` annotations allow configuration to be included depending on whether the application is a 'web application'. A web application is any application that is using a Spring `WebApplicationContext`, defines a session scope or has a `StandardServletEnvironment`.

SpEL expression conditions

The `@ConditionalOnExpression` annotation allows configuration to be included based on the result of a [SpEL expression](#).

46.4 Creating your own starter

A full Spring Boot starter for a library may contain the following components:

- The `autoconfigure` module that contains the auto-configuration code.
- The `starter` module that provides a dependency to the `autoconfigure` module as well as the library and any additional dependencies that are typically useful. In a nutshell, adding the starter should be enough to start using that library.

Tip

You may combine the auto-configuration code and the dependency management in a single module if you don't need to separate those two concerns.

Naming

Please make sure to provide a proper namespace for your starter. Do not start your module names with `spring-boot`, even if you are using a different Maven groupId. We may offer an official support for the thing you're auto-configuring in the future.

Here is a rule of thumb. Let's assume that you are creating a starter for "acme", name the auto-configure module `acme-spring-boot-autoconfigure` and the starter `acme-spring-boot-starter`. If you only have one module combining the two, use `acme-spring-boot-starter`.

Besides, if your starter provides configuration keys, use a proper namespace for them. In particular, do not include your keys in the namespaces that Spring Boot uses (e.g. `server`, `management`, `spring`,

etc). These are "ours" and we may improve/modify them in the future in such a way it could break your things.

Make sure to [trigger meta-data generation](#) so that IDE assistance is available for your keys as well. You may want to review the generated meta-data (`META-INF/spring-configuration-metadata.json`) to make sure your keys are properly documented.

Autoconfigure module

The autoconfigure module contains everything that is necessary to get started with the library. It may also contain configuration keys definition (`@ConfigurationProperties`) and any callback interface that can be used to further customize how the components are initialized.

Tip

You should mark the dependencies to the library as optional so that you can include the autoconfigure module in your projects more easily. If you do it that way, the library won't be provided and Spring Boot will back off by default.

Starter module

The starter is an empty jar, really. Its only purpose is to provide the necessary dependencies to work with the library; see it as an opinionated view of what is required to get started.

Do not make assumptions about the project in which your starter is added. If the library you are auto-configuring typically requires other starters, mention them as well. Providing a proper set of *default* dependencies may be hard if the number of optional dependencies is high as you should avoid bringing unnecessary dependencies for a typical usage of the library.

47. What to read next

If you want to learn more about any of the classes discussed in this section you can check out the [Spring Boot API documentation](#) or you can browse the [source code directly](#). If you have specific questions, take a look at the [how-to](#) section.

If you are comfortable with Spring Boot's core features, you can carry on and read about [production-ready features](#).

Part V. Spring Boot Actuator: Production-ready features

Spring Boot includes a number of additional features to help you monitor and manage your application when it's pushed to production. You can choose to manage and monitor your application using HTTP endpoints or with JMX. Auditing, health and metrics gathering can be automatically applied to your application.

Actuator HTTP endpoints are only available with a Spring MVC-based application. In particular, it will not work with Jersey [unless you enable Spring MVC as well.](#)

48. Enabling production-ready features

The `spring-boot-actuator` module provides all of Spring Boot's production-ready features. The simplest way to enable the features is to add a dependency to the `spring-boot-starter-actuator` 'Starter'.

Definition of Actuator

An actuator is a manufacturing term, referring to a mechanical device for moving or controlling something. Actuators can generate a large amount of motion from a small change.

To add the actuator to a Maven based project, add the following 'Starter' dependency:

```
<dependencies>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-actuator</artifactId>
  </dependency>
</dependencies>
```

For Gradle, use the declaration:

```
dependencies {
    compile("org.springframework.boot:spring-boot-starter-actuator")
}
```

49. Endpoints

Actuator endpoints allow you to monitor and interact with your application. Spring Boot includes a number of built-in endpoints and you can also add your own. For example the `health` endpoint provides basic application health information.

The way that endpoints are exposed will depend on the type of technology that you choose. Most applications choose HTTP monitoring, where the ID of the endpoint along with a prefix of `/application` is mapped to a URL. For example, by default, the `health` endpoint will be mapped to `/application/health`.

The following technology agnostic endpoints are available:

ID	Description	Sensitive Default
<code>actuator</code>	Provides a hypermedia-based “discovery page” for the other endpoints. Requires Spring HATEOAS to be on the classpath.	true
<code>auditevents</code>	Exposes audit events information for the current application.	true
<code>autoconfig</code>	Displays an auto-configuration report showing all auto-configuration candidates and the reason why they ‘were’ or ‘were not’ applied.	true
<code>beans</code>	Displays a complete list of all the Spring beans in your application.	true
<code>configprops</code>	Displays a collated list of all <code>@ConfigurationProperties</code> .	true
<code>dump</code>	Performs a thread dump.	true
<code>env</code>	Exposes properties from Spring’s <code>ConfigurableEnvironment</code> .	true
<code>flyway</code>	Shows any Flyway database migrations that have been applied.	true
<code>health</code>	Shows application health information (when the application is secure, a simple ‘status’ when accessed over an unauthenticated connection or full message details when authenticated).	false
<code>info</code>	Displays arbitrary application info.	false
<code>loggers</code>	Shows and modifies the configuration of loggers in the application.	true
<code>liquibase</code>	Shows any Liquibase database migrations that have been applied.	true
<code>metrics</code>	Shows ‘metrics’ information for the current application.	true
<code>mappings</code>	Displays a collated list of all <code>@RequestMapping</code> paths.	true

ID	Description	Sensitive Default
shutdown	Allows the application to be gracefully shutdown (not enabled by default).	true
trace	Displays trace information (by default the last 100 HTTP requests).	true

If you are using Spring MVC, the following additional endpoints can also be used:

ID	Description	Sensitive Default
docs	Displays documentation, including example requests and responses, for the Actuator's endpoints. Requires <code>spring-boot-actuator-docs</code> to be on the classpath.	false
heapdump	Returns a GZip compressed <code>hprof</code> heap dump file.	true
jolokia	Exposes JMX beans over HTTP (when Jolokia is on the classpath).	true
logfile	Returns the contents of the logfile (if <code>logging.file</code> or <code>logging.path</code> properties have been set). Supports the use of the HTTP <code>Range</code> header to retrieve part of the log file's content.	true

Note

Depending on how an endpoint is exposed, the `sensitive` property may be used as a security hint. For example, sensitive endpoints will require a username/password when they are accessed over HTTP (or simply disabled if web security is not enabled).

49.1 Customizing endpoints

Endpoints can be customized using Spring properties. You can change if an endpoint is `enabled`, if it is considered `sensitive` and even its `id`.

For example, here is an `application.properties` that changes the sensitivity and id of the beans endpoint and also enables `shutdown`.

```
endpoints.beans.id=springbeans
endpoints.beans.sensitive=false
endpoints.shutdown.enabled=true
```

Note

The prefix `#endpoints + . + name` is used to uniquely identify the endpoint that is being configured.

By default, all endpoints except for `shutdown` are enabled. If you prefer to specifically “opt-in” endpoint enablement you can use the `endpoints.enabled` property. For example, the following will disable *all* endpoints except for `info`:

```
endpoints.enabled=false
endpoints.info.enabled=true
```

Likewise, you can also choose to globally set the “sensitive” flag of all endpoints. By default, the sensitive flag depends on the type of endpoint (see the table above). For example, to mark *all* endpoints as sensitive except `info`:

```
endpoints.sensitive=true
endpoints.info.sensitive=false
```

49.2 Hypermedia for actuator MVC endpoints

If `endpoints.hypermedia.enabled` is set to `true` and [Spring HATEOAS](#) is on the classpath (e.g. through the `spring-boot-starter-hateoas` or if you are using [Spring Data REST](#)) then the HTTP endpoints from the Actuator are enhanced with hypermedia links, and a “discovery page” is added with links to all the endpoints. The “discovery page” is available on `/application` by default. It is implemented as an endpoint, allowing properties to be used to configure its path (`endpoints.actuator.path`) and whether or not it is enabled (`endpoints.actuator.enabled`).

When a custom management context path is configured, the “discovery page” will automatically move from `/application` to the root of the management context. For example, if the management context path is `/management` then the discovery page will be available from `/management`.

If the [HAL Browser](#) is on the classpath via its webjar (`org.webjars:hal-browser`), or via the `spring-data-rest-hal-browser` then an HTML “discovery page”, in the form of the HAL Browser, is also provided.

49.3 CORS support

[Cross-origin resource sharing](#) (CORS) is a [W3C specification](#) that allows you to specify in a flexible way what kind of cross domain requests are authorized. Actuator’s MVC endpoints can be configured to support such scenarios.

CORS support is disabled by default and is only enabled once the `endpoints.cors.allowed-origins` property has been set. The configuration below permits GET and POST calls from the `example.com` domain:

```
endpoints.cors.allowed-origins=http://example.com
endpoints.cors.allowed-methods=GET,POST
```

Tip

Check [EndpointCorsProperties](#) for a complete list of options.

49.4 Adding custom endpoints

If you add a `@Bean` of type `Endpoint` then it will automatically be exposed over JMX and HTTP (if there is a server available). An HTTP endpoints can be customized further by creating a bean of type `MvcEndpoint`. Your `MvcEndpoint` is not a `@Controller` but it can use `@RequestMapping` (and `@Managed*`) to expose resources.

Tip

If you are doing this as a library feature consider adding a configuration class annotated with `@ManagementContextConfiguration` to `/META-INF/spring.factories` under the key `org.springframework.boot.actuate.autoconfigure.ManagementContextConfiguration`. If you do that then the endpoint will move to a child context with all the other MVC endpoints if your users ask for a separate management port or address. A configuration declared this way can be a `WebConfigurerAdapter` if it wants to add static resources (for instance) to the management endpoints.

49.5 Health information

Health information can be used to check the status of your running application. It is often used by monitoring software to alert someone if a production system goes down. The default information exposed by the `health` endpoint depends on how it is accessed. For an unauthenticated connection in a secure application a simple 'status' message is returned, and for an authenticated connection additional details are also displayed (see [Section 50.7, "HTTP health endpoint format and access restrictions"](#) for HTTP details).

Health information is collected from all [HealthIndicator](#) beans defined in your `ApplicationContext`. Spring Boot includes a number of auto-configured `HealthIndicators` and you can also write your own. By default, the final system state is derived by the `HealthAggregator` which sorts the statuses from each `HealthIndicator` based on an ordered list of statuses. The first status in the sorted list is used as the overall health status. If no `HealthIndicator` returns a status that is known to the `HealthAggregator`, an UNKNOWN status is used.

49.6 Security with HealthIndicators

Information returned by `HealthIndicators` is often somewhat sensitive in nature. For example, you probably don't want to publish details of your database server to the world. For this reason, by default, only the health status is exposed over an unauthenticated HTTP connection. If you are happy for complete health information to always be exposed you can set `endpoints.health.sensitive` to `false`.

Health responses are also cached to prevent "denial of service" attacks. Use the `endpoints.health.time-to-live` property if you want to change the default cache period of 1000 milliseconds.

Auto-configured HealthIndicators

The following `HealthIndicators` are auto-configured by Spring Boot when appropriate:

Name	Description
CassandraHealthIndicator	Checks that a Cassandra database is up.
DiskSpaceHealthIndicator	Checks for low disk space.
DataSourceHealthIndicator	Checks that a connection to <code>DataSource</code> can be obtained.
ElasticsearchHealthIndicator	Checks that an Elasticsearch cluster is up.
JmsHealthIndicator	Checks that a JMS broker is up.

Name	Description
MailHealthIndicator	Checks that a mail server is up.
MongoHealthIndicator	Checks that a Mongo database is up.
RabbitHealthIndicator	Checks that a Rabbit server is up.
RedisHealthIndicator	Checks that a Redis server is up.
SolrHealthIndicator	Checks that a Solr server is up.

Tip

It is possible to disable them all using the `management.health.defaults.enabled` property.

Writing custom HealthIndicators

To provide custom health information you can register Spring beans that implement the [HealthIndicator](#) interface. You need to provide an implementation of the `health()` method and return a `Health` response. The `Health` response should include a status and can optionally include additional details to be displayed.

```
import org.springframework.boot.actuate.health.Health;
import org.springframework.boot.actuate.health.HealthIndicator;
import org.springframework.stereotype.Component;

@Component
public class MyHealthIndicator implements HealthIndicator {

    @Override
    public Health health() {
        int errorCode = check(); // perform some specific health check
        if (errorCode != 0) {
            return Health.down().withDetail("Error Code", errorCode).build();
        }
        return Health.up().build();
    }
}
```

Note

The identifier for a given `HealthIndicator` is the name of the bean without the `HealthIndicator` suffix if it exists. In the example above, the health information will be available in an entry named `my`.

In addition to Spring Boot's predefined [Status](#) types, it is also possible for `Health` to return a custom `Status` that represents a new system state. In such cases a custom implementation of the [HealthAggregator](#) interface also needs to be provided, or the default implementation has to be configured using the `management.health.status.order` configuration property.

For example, assuming a new `Status` with code `FATAL` is being used in one of your `HealthIndicator` implementations. To configure the severity order add the following to your application properties:

```
management.health.status.order=FATAL, DOWN, OUT_OF_SERVICE, UNKNOWN, UP
```

The HTTP status code in the response reflects the overall health status (e.g. UP maps to 200, OUT_OF_SERVICE or DOWN to 503). You might also want to register custom status mappings with the `HealthMvcEndpoint` if you access the health endpoint over HTTP. For example, the following maps FATAL to `HttpStatus.SERVICE_UNAVAILABLE`:

```
endpoints.health.mapping.FATAL=503
```

The default status mappings for the built-in statuses are:

Status	Mapping
DOWN	SERVICE_UNAVAILABLE (503)
OUT_OF_SERVICE	SERVICE_UNAVAILABLE (503)
UP	No mapping by default, so http status is 200
UNKNOWN	No mapping by default, so http status is 200

49.7 Application information

Application information exposes various information collected from all `InfoContributor` beans defined in your `ApplicationContext`. Spring Boot includes a number of auto-configured `InfoContributors` and you can also write your own.

Auto-configured InfoContributors

The following `InfoContributors` are auto-configured by Spring Boot when appropriate:

Name	Description
<code>EnvironmentInfoContributor</code>	Expose any key from the Environment under the info key.
<code>GitInfoContributor</code>	Expose git information if a <code>git.properties</code> file is available.
<code>BuildInfoContributor</code>	Expose build information if a <code>META-INF/build-info.properties</code> file is available.

Tip

It is possible to disable them all using the `management.info.defaults.enabled` property.

Custom application info information

You can customize the data exposed by the `info` endpoint by setting `info.*` Spring properties. All Environment properties under the info key will be automatically exposed. For example, you could add the following to your `application.properties`:

```
info.app.encoding=UTF-8
info.app.java.source=1.8
info.app.java.target=1.8
```

Tip

Rather than hardcoding those values you could also [expand info properties at build time](#).

Assuming you are using Maven, you could rewrite the example above as follows:

```
info.app.encoding=@project.build.sourceEncoding@
info.app.java.source=@java.version@
info.app.java.target=@java.version@
```

Git commit information

Another useful feature of the `info` endpoint is its ability to publish information about the state of your git source code repository when the project was built. If a `GitProperties` bean is available, the `git.branch`, `git.commit.id` and `git.commit.time` properties will be exposed.

Tip

A `GitProperties` bean is auto-configured if a `git.properties` file is available at the root of the classpath. See [Generate git information](#) for more details.

If you want to display the full git information (i.e. the full content of `git.properties`), use the `management.info.git.mode` property:

```
management.info.git.mode=full
```

Build information

The `info` endpoint can also publish information about your build if a `BuildProperties` bean is available. This happens if a `META-INF/build-info.properties` file is available in the classpath.

Tip

The Maven and Gradle plugins can both generate that file, see [Generate build information](#) for more details.

Writing custom InfoContributors

To provide custom application information you can register Spring beans that implement the [InfoContributor](#) interface.

The example below contributes an `example` entry with a single value:

```
import java.util.Collections;

import org.springframework.boot.actuate.info.Info;
import org.springframework.boot.actuate.info.InfoContributor;
import org.springframework.stereotype.Component;

@Component
public class ExampleInfoContributor implements InfoContributor {

    @Override
    public void contribute(Info.Builder builder) {
        builder.withDetail("example",
            Collections.singletonMap("key", "value"));
    }
}
```

If you hit the `info` endpoint you should see a response that contains the following additional entry:

```
{  
  "example": {  
    "key" : "value"  
  }  
}
```

50. Monitoring and management over HTTP

If you are developing a Spring MVC application, Spring Boot Actuator will auto-configure all enabled endpoints to be exposed over HTTP. The default convention is to use the `id` of the endpoint with a prefix of `/application` as the URL path. For example, `health` is exposed as `/application/health`.

50.1 Accessing sensitive endpoints

By default all sensitive HTTP endpoints are secured such that only users that have an `ACTUATOR` role may access them. Security is enforced using the standard `HttpServletRequest.isUserInRole` method.

Tip

Use the `management.security.roles` property if you want something different to `ACTUATOR`.

If you are deploying applications behind a firewall, you may prefer that all your actuator endpoints can be accessed without requiring authentication. You can do this by changing the `management.security.enabled` property:

`application.properties`.

```
management.security.enabled=false
```

Note

By default, actuator endpoints are exposed on the same port that serves regular HTTP traffic. Take care not to accidentally expose sensitive information if you change the `management.security.enabled` property.

If you're deploying applications publicly, you may want to add 'Spring Security' to handle user authentication. When 'Spring Security' is added, by default 'basic' authentication will be used with the username `user` and a generated password (which is printed on the console when the application starts).

Tip

Generated passwords are logged as the application starts. Search for 'Using default security password'.

You can use Spring properties to change the username and password and to change the security role(s) required to access the endpoints. For example, you might set the following in your `application.properties`:

```
security.user.name=admin
security.user.password=secret
management.security.roles=SUPERUSER
```

If your application has custom security configuration and you want all your actuator endpoints to be accessible without authentication, you need to explicitly configure that in your security configuration. Along with that, you need to change the `management.security.enabled` property to `false`.

If your custom security configuration secures your actuator endpoints, you also need to ensure that the authenticated user has the roles specified under `management.security.roles`.

Tip

If you don't have a use case for exposing basic health information to unauthenticated users, and you have secured the actuator endpoints with custom security, you can set `management.security.enabled` to `false`. This will inform Spring Boot to skip the additional role check.

50.2 Customizing the management endpoint paths

Sometimes it is useful to customize the prefix for the management endpoints. For example, your application might already use `/application` for another purpose. You can use the `management.context-path` property to change the prefix for your management endpoint:

```
management.context-path=/manage
```

The `application.properties` example above will change the endpoint from `/application/{id}` to `/manage/{id}` (e.g. `/manage/info`).

You can also change the “path” of an endpoint (using `endpoints.{name}.path`) which then changes the default resource path for the MVC endpoint. There is no validation on those values (so you can use anything that is legal in a URL path). For example, to change the location of the `/health` endpoint to `/ping/me` you can set `endpoints.health.path=/ping/me`.

Note

Even if an endpoint path is configured separately, it is still relative to the `management.context-path`.

Tip

If you provide a custom `MvcEndpoint` remember to include a settable `path` property, and default it to `/ {id}` if you want your code to behave like the standard MVC endpoints. (Take a look at the `HealthMvcEndpoint` to see how you might do that.) If your custom endpoint is an `Endpoint` (not an `MvcEndpoint`) then Spring Boot will take care of the path for you.

50.3 Customizing the management server port

Exposing management endpoints using the default HTTP port is a sensible choice for cloud based deployments. If, however, your application runs inside your own data center you may prefer to expose endpoints using a different HTTP port.

The `management.port` property can be used to change the HTTP port.

```
management.port=8081
```

Since your management port is often protected by a firewall, and not exposed to the public you might not need security on the management endpoints, even if your main application is secure. In that case you will have Spring Security on the classpath, and you can disable management security like this:

```
management.security.enabled=false
```

(If you don't have Spring Security on the classpath then there is no need to explicitly disable the management security in this way, and it might even break the application.)

50.4 Configuring management-specific SSL

When configured to use a custom port, the management server can also be configured with its own SSL using the various `management.ssl.*` properties. For example, this allows a management server to be available via HTTP while the main application uses HTTPS:

```
server.port=8443
server.ssl.enabled=true
server.ssl.key-store=classpath:store.jks
server.ssl.key-password=secret
management.port=8080
management.ssl.enabled=false
```

Alternatively, both the main server and the management server can use SSL but with different key stores:

```
server.port=8443
server.ssl.enabled=true
server.ssl.key-store=classpath:main.jks
server.ssl.key-password=secret
management.port=8080
management.ssl.enabled=true
management.ssl.key-store=classpath:management.jks
management.ssl.key-password=secret
```

50.5 Customizing the management server address

You can customize the address that the management endpoints are available on by setting the `management.address` property. This can be useful if you want to listen only on an internal or ops-facing network, or to only listen for connections from `localhost`.

Note

You can only listen on a different address if the port is different to the main server port.

Here is an example `application.properties` that will not allow remote management connections:

```
management.port=8081
management.address=127.0.0.1
```

50.6 Disabling HTTP endpoints

If you don't want to expose endpoints over HTTP you can set the management port to `-1`:

```
management.port=-1
```

50.7 HTTP health endpoint format and access restrictions

The information exposed by the health endpoint varies depending on whether or not it's accessed anonymously, and whether or not the enclosing application is secure. By default, when accessed anonymously in a secure application, any details about the server's health are hidden and the endpoint will simply indicate whether or not the server is up or down. Furthermore the response is cached for a configurable period to prevent the endpoint being used in a denial of service attack. The `endpoints.health.time-to-live` property is used to configure the caching period in milliseconds. It defaults to 1000, i.e. one second.

Sample summarized HTTP response (default for anonymous request):

```
$ curl -i localhost:8080/health
HTTP/1.1 200
X-Application-Context: application
Content-Type: application/vnd.spring-boot.actuator.v2+json;charset=UTF-8
Content-Length: 15

{"status":"UP"}
```

Sample summarized HTTP response for status "DOWN" (notice the 503 status code):

```
$ curl -i localhost:8080/health
HTTP/1.1 503
X-Application-Context: application
Content-Type: application/vnd.spring-boot.actuator.v2+json;charset=UTF-8
Content-Length: 17

{"status":"DOWN"}
```

Sample detailed HTTP response:

```
$ curl -i localhost:8080/health
HTTP/1.1 200 OK
X-Application-Context: application
Content-Type: application/vnd.spring-boot.actuator.v2+json;charset=UTF-8
Content-Length: 221

{
  "status" : "UP",
  "diskSpace" : {
    "status" : "UP",
    "total" : 63251804160,
    "free" : 31316164608,
    "threshold" : 10485760
  },
  "db" : {
    "status" : "UP",
    "database" : "H2",
    "hello" : 1
  }
}
```

The above-described restrictions can be enhanced, thereby allowing only authenticated users full access to the health endpoint in a secure application. To do so, set `endpoints.health.sensitive` to `true`. Here's a summary of behavior (with default `sensitive` flag value "false" indicated in bold):

<code>management.security.enabled</code>	<code>endpoints.health.sensitive</code>	Unauthenticated	Authenticated (with right role)
false	*	Full content	Full content
true	false	Status only	Full content
true	true	No content	Full content

51. Monitoring and management over JMX

Java Management Extensions (JMX) provide a standard mechanism to monitor and manage applications. By default Spring Boot will expose management endpoints as JMX MBeans under the `org.springframework.boot` domain.

51.1 Customizing MBean names

The name of the MBean is usually generated from the `id` of the endpoint. For example the `health` endpoint is exposed as `org.springframework.boot/Endpoint/healthEndpoint`.

If your application contains more than one Spring `ApplicationContext` you may find that names clash. To solve this problem you can set the `endpoints.jmx.unique-names` property to `true` so that MBean names are always unique.

You can also customize the JMX domain under which endpoints are exposed. Here is an example `application.properties`:

```
endpoints.jmx.domain=myapp
endpoints.jmx.unique-names=true
```

51.2 Disabling JMX endpoints

If you don't want to expose endpoints over JMX you can set the `endpoints.jmx.enabled` property to `false`:

```
endpoints.jmx.enabled=false
```

51.3 Using Jolokia for JMX over HTTP

Jolokia is a JMX-HTTP bridge giving an alternative method of accessing JMX beans. To use Jolokia, simply include a dependency to `org.jolokia:jolokia-core`. For example, using Maven you would add the following:

```
<dependency>
  <groupId>org.jolokia</groupId>
  <artifactId>jolokia-core</artifactId>
</dependency>
```

Jolokia can then be accessed using `/jolokia` on your management HTTP server.

Customizing Jolokia

Jolokia has a number of settings that you would traditionally configure using servlet parameters. With Spring Boot you can use your `application.properties`, simply prefix the parameter with `jolokia.config.:`

```
jolokia.config.debug=true
```

Disabling Jolokia

If you are using Jolokia but you don't want Spring Boot to configure it, simply set the `endpoints.jolokia.enabled` property to `false`:

```
endpoints.jolokia.enabled=false
```

52. Loggers

Spring Boot Actuator includes the ability to view and configure the log levels of your application at runtime. You can view either the entire list or an individual logger's configuration which is made up of both the explicitly configured logging level as well as the effective logging level given to it by the logging framework. These levels can be:

- TRACE
- DEBUG
- INFO
- WARN
- ERROR
- FATAL
- OFF
- null

with `null` indicating that there is no explicit configuration.

52.1 Configure a Logger

In order to configure a given logger, you `POST` a partial entity to the resource's URI:

```
{  
  "configuredLevel": "DEBUG"  
}
```

Tip

You can also pass a `null configuredLevel` to "reset" the specific level of the logger (and use the default configuration instead).

53. Metrics

Spring Boot Actuator includes a metrics service with ‘gauge’ and ‘counter’ support. A ‘gauge’ records a single value; and a ‘counter’ records a delta (an increment or decrement). Spring Boot Actuator also provides a [PublicMetrics](#) interface that you can implement to expose metrics that you cannot record via one of those two mechanisms. Look at [SystemPublicMetrics](#) for an example.

Metrics for all HTTP requests are automatically recorded, so if you hit the `metrics` endpoint you should see a response similar to this:

```
{
  "counter.status.200.root": 20,
  "counter.status.200.metrics": 3,
  "counter.status.200.star-star": 5,
  "counter.status.401.root": 4,
  "gauge.response.star-star": 6,
  "gauge.response.root": 2,
  "gauge.response.metrics": 3,
  "classes": 5808,
  "classes.loaded": 5808,
  "classes.unloaded": 0,
  "heap": 3728384,
  "heap.committed": 986624,
  "heap.init": 262144,
  "heap.used": 52765,
  "nonheap": 0,
  "nonheap.committed": 77568,
  "nonheap.init": 2496,
  "nonheap.used": 75826,
  "mem": 986624,
  "mem.free": 933858,
  "processors": 8,
  "threads": 15,
  "threads.daemon": 11,
  "threads.peak": 15,
  "threads.totalStarted": 42,
  "uptime": 494836,
  "instance.uptime": 489782,
  "datasource.primary.active": 5,
  "datasource.primary.usage": 0.25
}
```

Here we can see basic memory, heap, class loading, processor and thread pool information along with some HTTP metrics. In this instance the `root` (`/`) and `/metrics` URLs have returned HTTP 200 responses 20 and 3 times respectively. It also appears that the `root` URL returned HTTP 401 (unauthorized) 4 times. The double asterisks (`star-star`) comes from a request matched by Spring MVC as `/**` (normally a static resource).

The `gauge` shows the last response time for a request. So the last request to `root` took 2ms to respond and the last to `/metrics` took 3ms.

Note

In this example we are actually accessing the endpoint over HTTP using the `/metrics` URL, this explains why `metrics` appears in the response.

53.1 System metrics

The following system metrics are exposed by Spring Boot:

- The total system memory in KB (`mem`)
- The amount of free memory in KB (`mem.free`)
- The number of processors (`processors`)
- The system uptime in milliseconds (`uptime`)
- The application context uptime in milliseconds (`instance.uptime`)
- The average system load (`systemload.average`)
- Heap information in KB (`heap`, `heap.committed`, `heap.init`, `heap.used`)
- Thread information (`threads`, `thread.peak`, `thread.daemon`)
- Class load information (`classes`, `classes.loaded`, `classes.unloaded`)
- Garbage collection information (`gc.xxx.count`, `gc.xxx.time`)

53.2 DataSource metrics

The following metrics are exposed for each supported `DataSource` defined in your application:

- The number of active connections (`datasource.xxx.active`)
- The current usage of the connection pool (`datasource.xxx.usage`).

All data source metrics share the `datasource.` prefix. The prefix is further qualified for each data source:

- If the data source is the primary data source (that is either the only available data source or the one flagged `@Primary` amongst the existing ones), the prefix is `datasource.primary`.
- If the data source bean name ends with `DataSource`, the prefix is the name of the bean without `DataSource` (i.e. `datasource.batch` for `batchDataSource`).
- In all other cases, the name of the bean is used.

It is possible to override part or all of those defaults by registering a bean with a customized version of `DataSourcePublicMetrics`. By default, Spring Boot provides metadata for all supported data sources; you can add additional `DataSourcePoolMetadataProvider` beans if your favorite data source isn't supported out of the box. See `DataSourcePoolMetadataProvidersConfiguration` for examples.

53.3 Cache metrics

The following metrics are exposed for each supported cache defined in your application:

- The current size of the cache (`cache.xxx.size`)
- Hit ratio (`cache.xxx.hit.ratio`)
- Miss ratio (`cache.xxx.miss.ratio`)

Note

Cache providers do not expose the hit/miss ratio in a consistent way. While some expose an **aggregated** value (i.e. the hit ratio since the last time the stats were cleared), others expose a

temporal value (i.e. the hit ratio of the last second). Check your caching provider documentation for more details.

If two different cache managers happen to define the same cache, the name of the cache is prefixed by the name of the `CacheManager` bean.

It is possible to override part or all of those defaults by registering a bean with a customized version of `CachePublicMetrics`. By default, Spring Boot provides cache statistics for EhCache, Hazelcast, Infinispan, JCache and Caffeine. You can add additional `CacheStatisticsProvider` beans if your favorite caching library isn't supported out of the box. See `CacheStatisticsAutoConfiguration` for examples.

53.4 Tomcat session metrics

If you are using Tomcat as your embedded servlet container, session metrics will automatically be exposed. The `http.sessions.active` and `http.sessions.max` keys provide the number of active and maximum sessions.

53.5 Recording your own metrics

To record your own metrics inject a [CounterService](#) and/or [GaugeService](#) into your bean. The `CounterService` exposes `increment`, `decrement` and `reset` methods; the `GaugeService` provides a `submit` method.

Here is a simple example that counts the number of times that a method is invoked:

```
import org.springframework.beans.factory.annotation.Autowired;
import org.springframework.boot.actuate.metrics.CounterService;
import org.springframework.stereotype.Service;

@Service
public class MyService {

    private final CounterService counterService;

    @Autowired
    public MyService(CounterService counterService) {
        this.counterService = counterService;
    }

    public void exampleMethod() {
        this.counterService.increment("services.system.myService.invoked");
    }
}
```

Tip

You can use any string as a metric name but you should follow guidelines of your chosen store/graphing technology. Some good guidelines for Graphite are available on [Matt Aimonetti's Blog](#).

53.6 Adding your own public metrics

To add additional metrics that are computed every time the metrics endpoint is invoked, simply register additional `PublicMetrics` implementation bean(s). By default, all such beans are gathered by the endpoint. You can easily change that by defining your own `MetricsEndpoint`.

53.7 Metric writers, exporters and aggregation

Spring Boot provides a couple of implementations of a marker interface called `Exporter` which can be used to copy metric readings from the in-memory buffers to a place where they can be analyzed and displayed. Indeed, if you provide a `@Bean` that implements the `MetricWriter` interface (or `GaugeWriter` for simple use cases) and mark it `@ExportMetricWriter`, then it will automatically be hooked up to an `Exporter` and fed metric updates every 5 seconds (configured via `spring.metrics.export.delay-millis`). In addition, any `MetricReader` that you define and mark as `@ExportMetricReader` will have its values exported by the default exporter.

Note

This feature is enabling scheduling in your application (`@EnableScheduling`) which can be a problem if you run an integration test as your own scheduled tasks will start. You can disable this behaviour by setting `spring.metrics.export.enabled` to `false`.

The default exporter is a `MetricCopyExporter` which tries to optimize itself by not copying values that haven't changed since it was last called (the optimization can be switched off using a flag `spring.metrics.export.send-latest`). Note also that the Dropwizard `MetricRegistry` has no support for timestamps, so the optimization is not available if you are using Dropwizard metrics (all metrics will be copied on every tick).

The default values for the export trigger (`delay-millis`, `includes`, `excludes` and `send-latest`) can be set as `spring.metrics.export.*`. Individual values for specific `MetricWriters` can be set as `spring.metrics.export.triggers.<name>.*` where `<name>` is a bean name (or pattern for matching bean names).

Warning

The automatic export of metrics is disabled if you switch off the default `MetricRepository` (e.g. by using Dropwizard metrics). You can get back the same functionality by declaring a bean of your own of type `MetricReader` and declaring it to be `@ExportMetricReader`.

Example: Export to Redis

If you provide a `@Bean` of type `RedisMetricRepository` and mark it `@ExportMetricWriter` the metrics are exported to a Redis cache for aggregation. The `RedisMetricRepository` has two important parameters to configure it for this purpose: `prefix` and `key` (passed into its constructor). It is best to use a prefix that is unique to the application instance (e.g. using a random value and maybe the logical name of the application to make it possible to correlate with other instances of the same application). The “key” is used to keep a global index of all metric names, so it should be unique “globally”, whatever that means for your system (e.g. two instances of the same system could share a Redis cache if they have distinct keys).

Example:

```
@Bean
@ExportMetricWriter
MetricWriter metricWriter(MetricExportProperties export) {
    return new RedisMetricRepository(connectionFactory,
        export.getRedis().getPrefix(), export.getRedis().getKey());
}
```

application.properties.

```
spring.metrics.export.redis.prefix: metrics.mysystem.${spring.application.name:application}.
${random.value:0000}
spring.metrics.export.redis.key: keys.metrics.mysystem
```

The prefix is constructed with the application name and id at the end, so it can easily be used to identify a group of processes with the same logical name later.

Note

It's important to set both the `key` and the `prefix`. The key is used for all repository operations, and can be shared by multiple repositories. If multiple repositories share a key (like in the case where you need to aggregate across them), then you normally have a read-only “master” repository that has a short, but identifiable, prefix (like “metrics.mysystem”), and many write-only repositories with prefixes that start with the master prefix (like `metrics.mysystem.*` in the example above). It is efficient to read all the keys from a “master” repository like that, but inefficient to read a subset with a longer prefix (e.g. using one of the writing repositories).

Tip

The example above uses `MetricExportProperties` to inject and extract the key and prefix. This is provided to you as a convenience by Spring Boot, configured with sensible defaults. There is nothing to stop you using your own values as long as they follow the recommendations.

Example: Export to Open TSDB

If you provide a `@Bean` of type `OpenTsdB GaugeWriter` and mark it `@ExportMetricWriter` metrics are exported to [Open TSDB](#) for aggregation. The `OpenTsdB GaugeWriter` has a `url` property that you need to set to the Open TSDB “/put” endpoint, e.g. `localhost:4242/api/put`. It also has a `namingStrategy` that you can customize or configure to make the metrics match the data structure you need on the server. By default it just passes through the metric name as an Open TSDB metric name, and adds the tags “domain” (with value “org.springframework.metrics”) and “process” (with the value equal to the object hash of the naming strategy). Thus, after running the application and generating some metrics you can inspect the metrics in the TSD UI ([localhost:4242](#) by default).

Example:

```
curl localhost:4242/api/query?start=1h-ago&m=max:counter.status.200.root
[
  {
    "metric": "counter.status.200.root",
    "tags": {
      "domain": "org.springframework.metrics",
      "process": "b968a76"
    },
    "aggregateTags": [],
    "dps": {
      "1430492872": 2,
      "1430492875": 6
    }
  }
]
```

Example: Export to Statsd

To export metrics to Statsd, make sure first that you have added `com.timgroup:java-statsd-client` as a dependency of your project (Spring Boot provides a dependency management for it). Then add a `spring.metrics.export.statsd.host` value to your `application.properties` file. Connections will be opened to port 8125 unless a `spring.metrics.export.statsd.port` override is provided. You can use `spring.metrics.export.statsd.prefix` if you want a custom prefix.

Alternatively, you can provide a `@Bean` of type `StatsdMetricWriter` and mark it `@ExportMetricWriter`:

```
@Value("${spring.application.name:application}.${random.value:0000}")
private String prefix = "metrics";

@Bean
@ExportMetricWriter
MetricWriter metricWriter() {
    return new StatsdMetricWriter(prefix, "localhost", 8125);
}
```

Example: Export to JMX

If you provide a `@Bean` of type `JmxMetricWriter` marked `@ExportMetricWriter` the metrics are exported as MBeans to the local server (the `MBeanExporter` is provided by Spring Boot JMX auto-configuration as long as it is switched on). Metrics can then be inspected, graphed, alerted etc. using any tool that understands JMX (e.g. JConsole or JVisualVM).

Example:

```
@Bean
@ExportMetricWriter
MetricWriter metricWriter(MBeanExporter exporter) {
    return new JmxMetricWriter(exporter);
}
```

Each metric is exported as an individual MBean. The format for the `ObjectNames` is given by an `ObjectNamingStrategy` which can be injected into the `JmxMetricWriter` (the default breaks up the metric name and tags the first two period-separated sections in a way that should make the metrics group nicely in JVisualVM or JConsole).

53.8 Aggregating metrics from multiple sources

There is an `AggregateMetricReader` that you can use to consolidate metrics from different physical sources. Sources for the same logical metric just need to publish them with a period-separated prefix, and the reader will aggregate (by truncating the metric names, and dropping the prefix). Counters are summed and everything else (i.e. gauges) take their most recent value.

This is very useful if multiple application instances are feeding to a central (e.g. Redis) repository and you want to display the results. Particularly recommended in conjunction with a `MetricReaderPublicMetrics` for hooking up to the results to the `"/metrics"` endpoint.

Example:

```
@Autowired
private MetricExportProperties export;
```

```

@Bean
public PublicMetrics metricsAggregate() {
    return new MetricReaderPublicMetrics(aggregateMetricReader());
}

private MetricReader globalMetricsForAggregation() {
    return new RedisMetricRepository(this.connectionFactory,
        this.export.getRedis().getAggregatePrefix(), this.export.getRedis().getKey());
}

private MetricReader aggregateMetricReader() {
    AggregateMetricReader repository = new AggregateMetricReader(
        globalMetricsForAggregation());
    return repository;
}

```

Note

The example above uses `MetricExportProperties` to inject and extract the key and prefix. This is provided to you as a convenience by Spring Boot, and the defaults will be sensible. They are set up in `MetricExportAutoConfiguration`.

Note

The `MetricReaders` above are not `@Beans` and are not marked as `@ExportMetricReader` because they are just collecting and analyzing data from other repositories, and don't want to export their values.

53.9 Dropwizard Metrics

A default `MetricRegistry` Spring bean will be created when you declare a dependency to the `io.dropwizard.metrics:metrics-core` library; you can also register your own `@Bean` instance if you need customizations. Users of the [Dropwizard 'Metrics' library](#) will find that Spring Boot metrics are automatically published to `com.codahale.metrics.MetricRegistry`. Metrics from the `MetricRegistry` are also automatically exposed via the `/metrics` endpoint

When Dropwizard metrics are in use, the default `CounterService` and `GaugeService` are replaced with a `DropwizardMetricServices`, which is a wrapper around the `MetricRegistry` (so you can `@Autowired` one of those services and use it as normal). You can also create "special" Dropwizard metrics by prefixing your metric names with the appropriate type (i.e. `timer.*`, `histogram.*` for gauges, and `meter.*` for counters).

53.10 Message channel integration

If a `MessageChannel` bean called `metricsChannel` exists, then a `MetricWriter` will be created that writes metrics to that channel. Each message sent to the channel will contain a [Delta](#) or [Metric](#) payload and have a `metricName` header. The writer is automatically hooked up to an exporter (as for all writers), so all metric values will appear on the channel, and additional analysis or actions can be taken by subscribers (it's up to you to provide the channel and any subscribers you need).

54. Auditing

Spring Boot Actuator has a flexible audit framework that will publish events once Spring Security is in play ('authentication success', 'failure' and 'access denied' exceptions by default). This can be very useful for reporting, and also to implement a lock-out policy based on authentication failures. To customize published security events you can provide your own implementations of `AbstractAuthenticationAuditListener` and `AbstractAuthorizationAuditListener`.

You can also choose to use the audit services for your own business events. To do that you can either inject the existing `AuditEventRepository` into your own components and use that directly, or you can simply publish `AuditApplicationEvent` via the Spring `ApplicationEventPublisher` (using `ApplicationEventPublisherAware`).

55. Tracing

Tracing is automatically enabled for all HTTP requests. You can view the `trace` endpoint and obtain basic information about the last 100 requests:

```
[{
  "timestamp": 1394343677415,
  "info": {
    "method": "GET",
    "path": "/trace",
    "headers": {
      "request": {
        "Accept": "text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8",
        "Connection": "keep-alive",
        "Accept-Encoding": "gzip, deflate",
        "User-Agent": "Mozilla/5.0 Gecko/Firefox",
        "Accept-Language": "en-US,en;q=0.5",
        "Cookie": "_ga=GA1.1.827067509.1390890128; ..."
        "Authorization": "Basic ...",
        "Host": "localhost:8080"
      },
      "response": {
        "Strict-Transport-Security": "max-age=31536000 ; includeSubDomains",
        "X-Application-Context": "application:8080",
        "Content-Type": "application/json;charset=UTF-8",
        "status": "200"
      }
    }
  }
}, {
  "timestamp": 1394343684465,
  ...
}]
```

The following are included in the trace by default:

Name	Description
Request Headers	Headers from the request.
Response Headers	Headers from the response.
Cookies	Cookie from request headers and Set-Cookie from response headers.
Errors	The error attributes (if any).
Time Taken	The time taken to service the request in milliseconds.

55.1 Custom tracing

If you need to trace additional events you can inject a [TraceRepository](#) into your Spring beans. The `add` method accepts a single `Map` structure that will be converted to JSON and logged.

By default an `InMemoryTraceRepository` will be used that stores the last 100 events. You can define your own instance of the `InMemoryTraceRepository` bean if you need to expand the capacity. You can also create your own alternative `TraceRepository` implementation if needed.

56. Process monitoring

In Spring Boot Actuator you can find a couple of classes to create files that are useful for process monitoring:

- `ApplicationPidFileWriter` creates a file containing the application PID (by default in the application directory with the file name `application.pid`).
- `EmbeddedServerPortFileWriter` creates a file (or files) containing the ports of the embedded server (by default in the application directory with the file name `application.port`).

These writers are not activated by default, but you can enable them in one of the ways described below.

56.1 Extend configuration

In `META-INF/spring.factories` file you can activate the listener(s) that writes a PID file. Example:

```
org.springframework.context.ApplicationListener=\
org.springframework.boot.system.ApplicationPidFileWriter,\
org.springframework.boot.actuate.system.EmbeddedServerPortFileWriter
```

56.2 Programmatically

You can also activate a listener by invoking the `SpringApplication.addListener(...)` method and passing the appropriate `Writer` object. This method also allows you to customize the file name and path via the `Writer` constructor.

57. Cloud Foundry support

Spring Boot's actuator module includes additional support that is activated when you deploy to a compatible Cloud Foundry instance. The `/cloudfoundryapplication` path provides an alternative secured route to all `NamedMvcEndpoint` beans.

The extended support allows Cloud Foundry management UIs (such as the web application that you can use to view deployed applications) to be augmented with Spring Boot actuator information. For example, an application status page may include full health information instead of the typical "running" or "stopped" status.

Note

The `/cloudfoundryapplication` path is not directly accessible to regular users. In order to use the endpoint a valid UAA token must be passed with the request.

57.1 Disabling extended Cloud Foundry actuator support

If you want to fully disable the `/cloudfoundryapplication` endpoints you can add the following to your `application.properties` file:

application.properties.

```
management.cloudfoundry.enabled=false
```

57.2 Cloud Foundry self signed certificates

By default, the security verification for `/cloudfoundryapplication` endpoints makes SSL calls to various Cloud Foundry services. If your Cloud Foundry UAA or Cloud Controller services use self-signed certificates you will need to set the following property:

application.properties.

```
management.cloudfoundry.skip-ssl-validation=true
```

57.3 Custom security configuration

If you define custom security configuration, and you want extended Cloud Foundry actuator support, you'll should ensure that `/cloudfoundryapplication/**` paths are open. Without a direct open route, your Cloud Foundry application manager will not be able to obtain endpoint data.

For Spring Security, you'll typically include something like `mvcMatchers("/cloudfoundryapplication/**").permitAll()` in your configuration:

```
@Override
protected void configure(HttpSecurity http) throws Exception {
    http
        .authorizeRequests()
            .mvcMatchers("/cloudfoundryapplication/**")
                .permitAll()
            .mvcMatchers("/mypath")
                .hasAnyRole("SUPERUSER")
            .anyRequest()
                .authenticated().and()
            .httpBasic();
}
```

58. What to read next

If you want to explore some of the concepts discussed in this chapter, you can take a look at the actuator [sample applications](#). You also might want to read about graphing tools such as [Graphite](#).

Otherwise, you can continue on, to read about '[deployment options](#)' or jump ahead for some in-depth information about Spring Boot's [build tool plugins](#).

Part VI. Deploying Spring Boot applications

Spring Boot's flexible packaging options provide a great deal of choice when it comes to deploying your application. You can easily deploy Spring Boot applications to a variety of cloud platforms, to a container images (such as Docker) or to virtual/real machines.

This section covers some of the more common deployment scenarios.

59. Deploying to the cloud

Spring Boot's executable jars are ready-made for most popular cloud PaaS (platform-as-a-service) providers. These providers tend to require that you “bring your own container”; they manage application processes (not Java applications specifically), so they need some intermediary layer that adapts *your* application to the *cloud's* notion of a running process.

Two popular cloud providers, Heroku and Cloud Foundry, employ a “buildpack” approach. The buildpack wraps your deployed code in whatever is needed to *start* your application: it might be a JDK and a call to `java`, it might be an embedded web server, or it might be a full-fledged application server. A buildpack is pluggable, but ideally you should be able to get by with as few customizations to it as possible. This reduces the footprint of functionality that is not under your control. It minimizes divergence between development and production environments.

Ideally, your application, like a Spring Boot executable jar, has everything that it needs to run packaged within it.

In this section we'll look at what it takes to get the [simple application that we developed](#) in the “Getting Started” section up and running in the Cloud.

59.1 Cloud Foundry

Cloud Foundry provides default buildpacks that come into play if no other buildpack is specified. The Cloud Foundry [Java buildpack](#) has excellent support for Spring applications, including Spring Boot. You can deploy stand-alone executable jar applications, as well as traditional `.war` packaged applications.

Once you've built your application (using, for example, `mvn clean package`) and [installed the cf command line tool](#), simply deploy your application using the `cf push` command as follows, substituting the path to your compiled `.jar`. Be sure to have [logged in with your cf command line client](#) before pushing an application.

```
$ cf push acloudyspringtime -p target/demo-0.0.1-SNAPSHOT.jar
```

See the [cf push documentation](#) for more options. If there is a Cloud Foundry `manifest.yml` file present in the same directory, it will be consulted.

Note

Here we are substituting `acloudyspringtime` for whatever value you give `cf` as the name of your application.

At this point `cf` will start uploading your application:

```
Uploading acloudyspringtime... OK
Preparing to start acloudyspringtime... OK
-----> Downloaded app package (8.9M)
-----> Java Buildpack Version: v3.12 (offline) | https://github.com/cloudfoundry/java-buildpack.git#6f25b7e
-----> Downloading Open Jdk JRE 1.8.0_121 from https://java-buildpack.cloudfoundry.org/openjdk/trusty/x86_64/openjdk-1.8.0_121.tar.gz (found in cache)
      Expanding Open Jdk JRE to .java-buildpack/open_jdk_jre (1.6s)
-----> Downloading Open JDK Like Memory Calculator 2.0.2_RELEASE from https://java-buildpack.cloudfoundry.org/memory-calculator/trusty/x86_64/memory-calculator-2.0.2_RELEASE.tar.gz (found in cache)
```

```

Memory Settings: -Xss349K -Xmx681574K -XX:MaxMetaspaceSize=104857K -Xms681574K -
XX:MetaspaceSize=104857K
----> Downloading Container Certificate Trust Store 1.0.0_RELEASE from https://java-
buildpack.cloudfoundry.org/container-certificate-trust-store/container-certificate-trust-
store-1.0.0_RELEASE.jar (found in cache)
Adding certificates to .java-buildpack/container_certificate_trust_store/truststore.jks (0.6s)
----> Downloading Spring Auto Reconfiguration 1.10.0_RELEASE from https://java-
buildpack.cloudfoundry.org/auto-reconfiguration/auto-reconfiguration-1.10.0_RELEASE.jar (found in cache)
Checking status of app 'acloudyspringtime'...
 0 of 1 instances running (1 starting)
  ...
 0 of 1 instances running (1 starting)
  ...
 0 of 1 instances running (1 starting)
  ...
 1 of 1 instances running (1 running)

App started

```

Congratulations! The application is now live!

It's easy to then verify the status of the deployed application:

```

$ cf apps
Getting applications in ...
OK

name                requested state  instances  memory  disk  urls
...
acloudyspringtime  started          1/1        512M   1G    acloudyspringtime.cfapps.io
...

```

Once Cloud Foundry acknowledges that your application has been deployed, you should be able to hit the application at the URI given, in this case `http://acloudyspringtime.cfapps.io/`.

Binding to services

By default, metadata about the running application as well as service connection information is exposed to the application as environment variables (for example: `$VCAP_SERVICES`). This architecture decision is due to Cloud Foundry's polyglot (any language and platform can be supported as a buildpack) nature; process-scoped environment variables are language agnostic.

Environment variables don't always make for the easiest API so Spring Boot automatically extracts them and flattens the data into properties that can be accessed through Spring's `Environment` abstraction:

```

@Component
class MyBean implements EnvironmentAware {

    private String instanceId;

    @Override
    public void setEnvironment(Environment environment) {
        this.instanceId = environment.getProperty("vcap.application.instance_id");
    }

    // ...
}

```

All Cloud Foundry properties are prefixed with `vcap`. You can use `vcap` properties to access application information (such as the public URL of the application) and service information (such as database credentials). See `CloudFoundryVcapEnvironmentPostProcessor` Javadoc for complete details.

Tip

The [Spring Cloud Connectors](#) project is a better fit for tasks such as configuring a `DataSource`. Spring Boot includes auto-configuration support and a `spring-boot-starter-cloud-connectors` starter.

59.2 Heroku

Heroku is another popular PaaS platform. To customize Heroku builds, you provide a `Procfile`, which provides the incantation required to deploy an application. Heroku assigns a `port` for the Java application to use and then ensures that routing to the external URI works.

You must configure your application to listen on the correct port. Here's the `Procfile` for our starter REST application:

```
web: java -Dserver.port=$PORT -jar target/demo-0.0.1-SNAPSHOT.jar
```

Spring Boot makes `-D` arguments available as properties accessible from a Spring Environment instance. The `server.port` configuration property is fed to the embedded Tomcat, Jetty or Undertow instance which then uses it when it starts up. The `$PORT` environment variable is assigned to us by the Heroku PaaS.

This should be everything you need. The most common workflow for Heroku deployments is to `git push` the code to production.

```
$ git push heroku master

Initializing repository, done.
Counting objects: 95, done.
Delta compression using up to 8 threads.
Compressing objects: 100% (78/78), done.
Writing objects: 100% (95/95), 8.66 MiB | 606.00 KiB/s, done.
Total 95 (delta 31), reused 0 (delta 0)

-----> Java app detected
-----> Installing OpenJDK 1.8... done
-----> Installing Maven 3.3.1... done
-----> Installing settings.xml... done
-----> Executing: mvn -B -DskipTests=true clean install

[INFO] Scanning for projects...
Downloading: http://repo.spring.io/...
Downloaded: http://repo.spring.io/... (818 B at 1.8 KB/sec)
....
Downloaded: http://s3pository.heroku.com/jvm/... (152 KB at 595.3 KB/sec)
[INFO] Installing /tmp/build_0c35a5d2-a067-4abc-a232-14b1fb7a8229/target/...
[INFO] Installing /tmp/build_0c35a5d2-a067-4abc-a232-14b1fb7a8229/pom.xml ...
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----
[INFO] Total time: 59.358s
[INFO] Finished at: Fri Mar 07 07:28:25 UTC 2014
[INFO] Final Memory: 20M/493M
[INFO] -----

-----> Discovering process types
Procfile declares types -> web

-----> Compressing... done, 70.4MB
-----> Launching... done, v6
      http://agile-sierra-1405.herokuapp.com/ deployed to Heroku

To git@heroku.com:agile-sierra-1405.git
```

```
* [new branch]      master -> master
```

Your application should now be up and running on Heroku.

59.3 OpenShift

[OpenShift](#) is the RedHat public (and enterprise) PaaS solution. Like Heroku, it works by running scripts triggered by git commits, so you can script the launching of a Spring Boot application in pretty much any way you like as long as the Java runtime is available (which is a standard feature you can ask for at OpenShift). To do this you can use the [DIY Cartridge](#) and hooks in your repository under `.openshift/action_hooks`:

The basic model is to:

1. Ensure Java and your build tool are installed remotely, e.g. using a `pre_build` hook (Java and Maven are installed by default, Gradle is not)
2. Use a `build` hook to build your jar (using Maven or Gradle), e.g.

```
#!/bin/bash
cd $OPENSIFT_REPO_DIR
mvn package -s .openshift/settings.xml -DskipTests=true
```

3. Add a `start` hook that calls `java -jar ...`

```
#!/bin/bash
cd $OPENSIFT_REPO_DIR
nohup java -jar target/*.jar --server.port=${OPENSIFT_DIY_PORT} --server.address=${OPENSIFT_DIY_IP}
&
```

4. Use a `stop` hook (since the start is supposed to return cleanly), e.g.

```
#!/bin/bash
source $OPENSIFT_CARTRIDGE_SDK_BASH
PID=$(ps -ef | grep java.*\.jar | grep -v grep | awk '{ print $2 }')
if [ -z "$PID" ]
then
    client_result "Application is already stopped"
else
    kill $PID
fi
```

5. Embed service bindings from environment variables provided by the platform in your `application.properties`, e.g.

```
spring.datasource.url: jdbc:mysql://${OPENSIFT_MYSQL_DB_HOST}:${OPENSIFT_MYSQL_DB_PORT}/
${OPENSIFT_APP_NAME}
spring.datasource.username: ${OPENSIFT_MYSQL_DB_USERNAME}
spring.datasource.password: ${OPENSIFT_MYSQL_DB_PASSWORD}
```

There's a blog on [running Gradle in OpenShift](#) on their website that will get you started with a gradle build to run the app.

59.4 Amazon Web Services (AWS)

Amazon Web Services offers multiple ways to install Spring Boot based applications, either as traditional web applications (war) or as executable jar files with an embedded web server. Options include :

- AWS Elastic Beanstalk

- AWS Code Deploy
- AWS OPS Works
- AWS Cloud Formation
- AWS Container Registry

Each has different features and pricing model, here we will describe only the simplest option : AWS Elastic Beanstalk.

AWS Elastic Beanstalk

As described in the official [Elastic Beanstalk Java guide](#), there are two main options to deploy a Java application; You can either use the “Tomcat Platform” or the “Java SE platform”.

Using the Tomcat platform

This option applies to Spring Boot projects producing a war file. There is no any special configuration required, just follow the official guide.

Using the Java SE platform

This option applies to Spring Boot projects producing a jar file and running an embedded web container. Elastic Beanstalk environments run an nginx instance on port 80 to proxy the actual application, running on port 5000. To configure it, add the following to your `application.properties`:

```
server.port=5000
```

Best practices

Uploading binaries instead of sources

By default Elastic Beanstalk uploads sources and compiles them in AWS. To upload the binaries instead, add the following to your `.elasticbeanstalk/config.yml` file:

```
deploy:  
  artifact: target/demo-0.0.1-SNAPSHOT.jar
```

Reduce costs by setting the environment type

By default an Elastic Beanstalk environment is load balanced. The load balancer has a cost perspective, to avoid it, set the environment type to “Single instance” as described [in the Amazon documentation](#). Single instance environments can be created using the CLI as well using the following command:

```
eb create -s
```

Summary

This is one of the easiest ways to get to AWS, but there are more things to cover, e.g.: how to integrate Elastic Beanstalk into any CI / CD tool, using the Elastic Beanstalk maven plugin instead of the CLI, etc. There is a [blog](#) covering these topics more in detail.

59.5 Boxfuse and Amazon Web Services

[Boxfuse](#) works by turning your Spring Boot executable jar or war into a minimal VM image that can be deployed unchanged either on VirtualBox or on AWS. Boxfuse comes with deep integration for Spring

Boot and will use the information from your Spring Boot configuration file to automatically configure ports and health check URLs. Boxfuse leverages this information both for the images it produces as well as for all the resources it provisions (instances, security groups, elastic load balancers, etc).

Once you have created a [Boxfuse account](#), connected it to your AWS account, and installed the latest version of the Boxfuse Client, you can deploy your Spring Boot application to AWS as follows (ensure the application has been built by Maven or Gradle first using, for example, `mvn clean package`):

```
$ boxfuse run myapp-1.0.jar -env=prod
```

See the [boxfuse run documentation](#) for more options. If there is a [boxfuse.com/docs/commandline/#configuration](#) [boxfuse.conf] file present in the current directory, it will be consulted.

Tip

By default Boxfuse will activate a Spring profile named `boxfuse` on startup and if your executable jar or war contains an [boxfuse.com/docs/payloads/springboot.html#configuration](#) [application-boxfuse.properties] file, Boxfuse will base its configuration based on the properties it contains.

At this point `boxfuse` will create an image for your application, upload it, and then configure and start the necessary resources on AWS:

```
Fusing Image for myapp-1.0.jar ...
Image fused in 00:06.838s (53937 K) -> axelfontaine/myapp:1.0
Creating axelfontaine/myapp ...
Pushing axelfontaine/myapp:1.0 ...
Verifying axelfontaine/myapp:1.0 ...
Creating Elastic IP ...
Mapping myapp-axelfontaine.boxfuse.io to 52.28.233.167 ...
Waiting for AWS to create an AMI for axelfontaine/myapp:1.0 in eu-central-1 (this may take up to 50
seconds) ...
AMI created in 00:23.557s -> ami-d23f38cf
Creating security group boxfuse-sg_axelfontaine/myapp:1.0 ...
Launching t2.micro instance of axelfontaine/myapp:1.0 (ami-d23f38cf) in eu-central-1 ...
Instance launched in 00:30.306s -> i-92ef9f53
Waiting for AWS to boot Instance i-92ef9f53 and Payload to start at http://52.28.235.61/ ...
Payload started in 00:29.266s -> http://52.28.235.61/
Remapping Elastic IP 52.28.233.167 to i-92ef9f53 ...
Waiting 15s for AWS to complete Elastic IP Zero Downtime transition ...
Deployment completed successfully. axelfontaine/myapp:1.0 is up and running at http://myapp-
axelfontaine.boxfuse.io/
```

Your application should now be up and running on AWS.

There's a blog on [deploying Spring Boot apps on EC2](#) as well as [documentation for the Boxfuse Spring Boot integration](#) on their website that will get you started with a Maven build to run the app.

59.6 Google Cloud

Google Cloud has several options that could be used to launch Spring Boot applications. The easiest to get started with is probably App Engine, but you could also find ways to run Spring Boot in a container with Container Engine, or on a virtual machine using Compute Engine.

To run in App Engine you can create a project in the UI first, which sets up a unique identifier for you and also HTTP routes. Add a Java app to the project and leave it empty, then use the [Google Cloud SDK](#) to push your Spring Boot app into that slot from the command line or CI build.

App Engine needs you to create an `app.yaml` file to describe the resources your app requires. Normally you put this in `src/main/appengine`, and it looks something like this:

```
service: default

runtime: java
env: flex

runtime_config:
  jdk: openjdk8

handlers:
- url: /*
  script: this field is required, but ignored

manual_scaling:
  instances: 1

health_check:
  enable_health_check: False

env_variables:
  ENCRYPT_KEY: your_encryption_key_here
```

You can deploy the app, for example, with a Maven plugin by simply adding the project ID to the build configuration:

```
<plugin>
  <groupId>com.google.cloud.tools</groupId>
  <artifactId>appengine-maven-plugin</artifactId>
  <version>1.3.0</version>
  <configuration>
    <project>myproject</project>
  </configuration>
</plugin>
```

Then deploy with `mvn appengine:deploy` (if you need to authenticate first the build will fail).

Note

Google App Engine Classic is tied to the Servlet 2.5 API, so you can't deploy a Spring Application there without some modifications. See the [Servlet 2.5 section](#) of this guide.

60. Installing Spring Boot applications

In addition to running Spring Boot applications using `java -jar` it is also possible to make fully executable applications for Unix systems. A fully executable jar can be executed like any other executable binary or it can be [registered with `init.d` or `systemd`](#). This makes it very easy to install and manage Spring Boot applications in common production environments.

Warning

Fully executable jars work by embedding an extra script at the front of the file. Currently, some tools do not accept this format so you may not always be able to use this technique. For example, `jar -xf` may silently fail to extract a jar or war that has been made fully-executable. It is recommended that you only make your jar or war fully executable if you intend to execute it directly, rather than running it with `java -jar` or deploying it to a servlet container.

To create a ‘fully executable’ jar with Maven use the following plugin configuration:

```
<plugin>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-maven-plugin</artifactId>
  <configuration>
    <executable>true</executable>
  </configuration>
</plugin>
```

With Gradle, the equivalent configuration is:

```
springBoot {
    executable = true
}
```

You can then run your application by typing `./my-application.jar` (where `my-application` is the name of your artifact). The directory containing the jar will be used as your application’s working directory.

60.1 Supported operating systems

The default script supports most Linux distributions and is tested on CentOS and Ubuntu. Other platforms, such as OS X and FreeBSD, will require the use of a custom `embeddedLaunchScript`.

60.2 Unix/Linux services

Spring Boot application can be easily started as Unix/Linux services using either `init.d` or `systemd`.

Installation as an `init.d` service (System V)

If you’ve configured Spring Boot’s Maven or Gradle plugin to generate a [fully executable jar](#), and you’re not using a custom `embeddedLaunchScript`, then your application can be used as an `init.d` service. Simply symlink the jar to `init.d` to support the standard `start`, `stop`, `restart` and `status` commands.

The script supports the following features:

- Starts the services as the user that owns the jar file
- Tracks application’s PID using `/var/run/<appname>/<appname>.pid`

- Writes console logs to `/var/log/<appname>.log`

Assuming that you have a Spring Boot application installed in `/var/myapp`, to install a Spring Boot application as an `init.d` service simply create a symlink:

```
$ sudo ln -s /var/myapp/myapp.jar /etc/init.d/myapp
```

Once installed, you can start and stop the service in the usual way. For example, on a Debian based system:

```
$ service myapp start
```

Tip

If your application fails to start, check the log file written to `/var/log/<appname>.log` for errors.

You can also flag the application to start automatically using your standard operating system tools. For example, on Debian:

```
$ update-rc.d myapp defaults <priority>
```

Securing an init.d service

Note

The following is a set of guidelines on how to secure a Spring Boot application that's being run as an `init.d` service. It is not intended to be an exhaustive list of everything that should be done to harden an application and the environment in which it runs.

When executed as `root`, as is the case when `root` is being used to start an `init.d` service, the default executable script will run the application as the user which owns the `jar` file. You should never run a Spring Boot application as `root` so your application's `jar` file should never be owned by `root`. Instead, create a specific user to run your application and use `chown` to make it the owner of the `jar` file. For example:

```
$ chown bootapp:bootapp your-app.jar
```

In this case, the default executable script will run the application as the `bootapp` user.

Tip

To reduce the chances of the application's user account being compromised, you should consider preventing it from using a login shell. Set the account's shell to `/usr/sbin/nologin`, for example.

You should also take steps to prevent the modification of your application's `jar` file. Firstly, configure its permissions so that it cannot be written and can only be read or executed by its owner:

```
$ chmod 500 your-app.jar
```

Secondly, you should also take steps to limit the damage if your application or the account that's running it is compromised. If an attacker does gain access, they could make the `jar` file writable and change its contents. One way to protect against this is to make it immutable using `chattr`:

```
$ sudo chmod +x your-app.jar
```

This will prevent any user, including root, from modifying the jar.

If root is used to control the application's service and you [use a .conf file](#) to customize its startup, the .conf file will be read and evaluated by the root user. It should be secured accordingly. Use `chmod` so that the file can only be read by the owner and use `chown` to make root the owner:

```
$ chmod 400 your-app.conf
$ sudo chown root:root your-app.conf
```

Installation as a systemd service

Systemd is the successor of the System V init system, and is now being used by many modern Linux distributions. Although you can continue to use `init.d` scripts with `systemd`, it is also possible to launch Spring Boot applications using `systemd` 'service' scripts.

Assuming that you have a Spring Boot application installed in `/var/myapp`, to install a Spring Boot application as a `systemd` service create a script named `myapp.service` using the following example and place it in `/etc/systemd/system` directory:

```
[Unit]
Description=myapp
After=syslog.target

[Service]
User=myapp
ExecStart=/var/myapp/myapp.jar
SuccessExitStatus=143

[Install]
WantedBy=multi-user.target
```

Tip

Remember to change the `Description`, `User` and `ExecStart` fields for your application.

Tip

Note that `ExecStart` field does not declare the script action command, which means that `run` command is used by default.

Note that unlike when running as an `init.d` service, user that runs the application, PID file and console log file are managed by `systemd` itself and therefore must be configured using appropriate fields in 'service' script. Consult the [service unit configuration man page](#) for more details.

To flag the application to start automatically on system boot use the following command:

```
$ systemctl enable myapp.service
```

Refer to `man systemctl` for more details.

Customizing the startup script

The default embedded startup script written by the Maven or Gradle plugin can be customized in a number of ways. For most people, using the default script along with a few customizations is

usually enough. If you find you can't customize something that you need to, you can always use the `embeddedLaunchScript` option to write your own file entirely.

Customizing script when it's written

It often makes sense to customize elements of the start script as it's written into the jar file. For example, `init.d` scripts can provide a "description" and, since you know this up front (and it won't change), you may as well provide it when the jar is generated.

To customize written elements, use the `embeddedLaunchScriptProperties` option of the Spring Boot Maven or Gradle plugins.

The following property substitutions are supported with the default script:

Name	Description
<code>mode</code>	The script mode. Defaults to <code>auto</code> .
<code>initInfoProjectName</code>	The <code>name</code> section of "INIT INFO". Defaults to <code>spring-boot-application</code> for Gradle and to <code>\${project.artifactId}</code> for Maven.
<code>initInfoRequiredStart</code>	The <code>required-start</code> section of "INIT INFO". Defaults to <code>\$remote_fs \$syslog \$network</code> .
<code>initInfoRequiredStop</code>	The <code>required-stop</code> section of "INIT INFO". Defaults to <code>\$remote_fs \$syslog \$network</code> .
<code>initInfoDefaultStart</code>	The <code>default-start</code> section of "INIT INFO". Defaults to <code>2 3 4 5</code> .
<code>initInfoDefaultStop</code>	The <code>default-stop</code> section of "INIT INFO". Defaults to <code>0 1 6</code> .
<code>initInfoShortDescription</code>	The <code>short-description</code> section of "INIT INFO". Defaults to <code>Spring Boot Application for Gradle</code> and to <code>\${project.name}</code> for Maven.
<code>initInfoDescription</code>	The <code>description</code> section of "INIT INFO". Defaults to <code>Spring Boot Application for Gradle</code> and to <code>\${project.description}</code> (falling back to <code>\${project.name}</code>) for Maven.
<code>initInfoChkconfig</code>	The <code>chkconfig</code> section of "INIT INFO". Defaults to <code>2345 99 01</code> .
<code>confFolder</code>	The default value for <code>CONF_FOLDER</code> . Defaults to the folder containing the jar.
<code>logFolder</code>	The default value for <code>LOG_FOLDER</code> . Only valid for an <code>init.d</code> service.
<code>logFilename</code>	The default value for <code>LOG_FILENAME</code> . Only valid for an <code>init.d</code> service.
<code>pidFolder</code>	The default value for <code>PID_FOLDER</code> . Only valid for an <code>init.d</code> service.
<code>pidFilename</code>	The default value for the name of the pid file in <code>PID_FOLDER</code> . Only valid for an <code>init.d</code> service.
<code>useStartStopDaemon</code>	If the <code>start-stop-daemon</code> command, when it's available, should be used to control the process. Defaults to <code>true</code> .
<code>stopWaitTime</code>	The default value for <code>STOP_WAIT_TIME</code> . Only valid for an <code>init.d</code> service. Defaults to 60 seconds.

Customizing script when it runs

For items of the script that need to be customized *after* the jar has been written you can use environment variables or a [config file](#).

The following environment properties are supported with the default script:

Variable	Description
MODE	The “mode” of operation. The default depends on the way the jar was built, but will usually be <code>auto</code> (<i>meaning it tries to guess if it is an init script by checking if it is a symlink in a directory called <code>init.d</code></i>). You can explicitly set it to <code>service</code> so that the <code>stop start status restart</code> commands work, or to <code>run</code> if you just want to run the script in the foreground.
USE_START_SCRIPT	If the <code>systemd</code> <code>stop-daemon</code> command, when it's available, should be used to control the process. Defaults to <code>true</code> .
PID_FOLDER	The root name of the pid folder (<code>/var/run</code> by default).
LOG_FOLDER	The name of the folder to put log files in (<code>/var/log</code> by default).
CONF_FOLDER	The name of the folder to read <code>.conf</code> files from (same folder as <code>jar-file</code> by default).
LOG_FILENAME	The name of the log file in the <code>LOG_FOLDER</code> (<code><appname>.log</code> by default).
APP_NAME	The name of the app. If the jar is run from a symlink the script guesses the app name, but if it is not a symlink, or you want to explicitly set the app name this can be useful.
RUN_ARGS	The arguments to pass to the program (the Spring Boot app).
JAVA_HOME	The location of the <code>java</code> executable is discovered by using the <code>PATH</code> by default, but you can set it explicitly if there is an executable file at <code>\$JAVA_HOME/bin/java</code> .
JAVA_OPTS	Options that are passed to the JVM when it is launched.
JARFILE	The explicit location of the jar file, in case the script is being used to launch a jar that it is not actually embedded in.
DEBUG	if not empty will set the <code>-x</code> flag on the shell process, making it easy to see the logic in the script.
STOP_WAIT_TIME	The time in seconds to wait when stopping the application before forcing a shutdown (<code>60</code> by default).

Note

The `PID_FOLDER`, `LOG_FOLDER` and `LOG_FILENAME` variables are only valid for an `init.d` service. With `systemd` the equivalent customizations are made using ‘service’ script. Check the [service unit configuration man page](#) for more details.

With the exception of `JARFILE` and `APP_NAME`, the above settings can be configured using a `.conf` file. The file is expected next to the jar file and have the same name but suffixed with `.conf` rather than `.jar`. For example, a jar named `/var/myapp/myapp.jar` will use the configuration file named `/var/myapp/myapp.conf`.

myapp.conf.

```
JAVA_OPTS=-Xmx1024M
LOG_FOLDER=/custom/log/folder
```

Tip

You can use a `CONF_FOLDER` environment variable to customize the location of the config file if you don't like it living next to the jar.

To learn about securing this file appropriately, please refer to [the guidelines for securing an init.d service](#).

60.3 Microsoft Windows services

Spring Boot application can be started as Windows service using [winsw](#).

A sample [maintained separately](#) to the core of Spring Boot describes step-by-step how you can create a Windows service for your Spring Boot application.

61. What to read next

Check out the [Cloud Foundry](#), [Heroku](#), [OpenShift](#) and [Boxfuse](#) web sites for more information about the kinds of features that a PaaS can offer. These are just four of the most popular Java PaaS providers, since Spring Boot is so amenable to cloud-based deployment you're free to consider other providers as well.

The next section goes on to cover the [Spring Boot CLI](#); or you can jump ahead to read about [build tool plugins](#).

Part VII. Spring Boot CLI

The Spring Boot CLI is a command line tool that can be used if you want to quickly develop with Spring. It allows you to run Groovy scripts, which means that you have a familiar Java-like syntax, without so much boilerplate code. You can also bootstrap a new project or write your own command for it.

62. Installing the CLI

The Spring Boot CLI can be installed manually; using SDKMAN! (the SDK Manager) or using Homebrew or MacPorts if you are an OSX user. See [Section 10.2, “Installing the Spring Boot CLI”](#) in the “Getting started” section for comprehensive installation instructions.

63. Using the CLI

Once you have installed the CLI you can run it by typing `spring`. If you run `spring` without any arguments, a simple help screen is displayed:

```
$ spring
usage: spring [--help] [--version]
       <command> [<args>]

Available commands are:

  run [options] <files> [--] [args]
    Run a spring groovy script

... more command help is shown here
```

You can use `help` to get more details about any of the supported commands. For example:

```
$ spring help run
spring run - Run a spring groovy script

usage: spring run [options] <files> [--] [args]

Option                Description
-----                -
--autoconfigure [Boolean] Add autoconfigure compiler
                        transformations (default: true)
--classpath, -cp       Additional classpath entries
-e, --edit              Open the file with the default system
                        editor
--no-guess-dependencies Do not attempt to guess dependencies
--no-guess-imports     Do not attempt to guess imports
-q, --quiet             Quiet logging
-v, --verbose           Verbose logging of dependency
                        resolution
--watch                Watch the specified file for changes
```

The `version` command provides a quick way to check which version of Spring Boot you are using.

```
$ spring version
Spring CLI v2.0.0.M3
```

63.1 Running applications using the CLI

You can compile and run Groovy source code using the `run` command. The Spring Boot CLI is completely self-contained so you don't need any external Groovy installation.

Here is an example “hello world” web application written in Groovy:

hello.groovy.

```
@RestController
class WebApplication {

    @RequestMapping("/")
    String home() {
        "Hello World!"
    }
}
```

To compile and run the application type:

```
$ spring run hello.groovy
```

To pass command line arguments to the application, you need to use a `--` to separate them from the “spring” command arguments, e.g.

```
$ spring run hello.groovy -- --server.port=9000
```

To set JVM command line arguments you can use the `JAVA_OPTS` environment variable, e.g.

```
$ JAVA_OPTS=-Xmx1024m spring run hello.groovy
```

Deduced “grab” dependencies

Standard Groovy includes a `@Grab` annotation which allows you to declare dependencies on a third-party libraries. This useful technique allows Groovy to download jars in the same way as Maven or Gradle would, but without requiring you to use a build tool.

Spring Boot extends this technique further, and will attempt to deduce which libraries to “grab” based on your code. For example, since the `WebApplication` code above uses `@RestController` annotations, “Tomcat” and “Spring MVC” will be grabbed.

The following items are used as “grab hints”:

Items	Grabs
JdbcTemplate, NamedParameterJdbcTemplate, DataSource	JDBC Application.
@EnableJms	JMS Application.
@EnableCaching	Caching abstraction.
@Test	JUnit.
@EnableRabbit	RabbitMQ.
@EnableReactor	Project Reactor.
extends Specification	Spock test.
@EnableBatchProcessing	Spring Batch.
@MessageEndpoint @EnableIntegrationPatterns	Spring Integration.
@EnableDeviceResolver	Spring Mobile.
@Controller @RestController @EnableWebMvc	Spring MVC + Embedded Tomcat.
@EnableWebSecurity	Spring Security.
@EnableTransactionManagement	Spring Transaction Management.

Tip

See subclasses of [CompilerAutoConfiguration](#) in the Spring Boot CLI source code to understand exactly how customizations are applied.

Deduced “grab” coordinates

Spring Boot extends Groovy’s standard `@Grab` support by allowing you to specify a dependency without a group or version, for example `@Grab('freemarker')`. This will consult Spring Boot’s default dependency metadata to deduce the artifact’s group and version. Note that the default metadata is tied to the version of the CLI that you’re using – it will only change when you move to a new version of the CLI, putting you in control of when the versions of your dependencies may change. A table showing the dependencies and their versions that are included in the default metadata can be found in the [appendix](#).

Default import statements

To help reduce the size of your Groovy code, several `import` statements are automatically included. Notice how the example above refers to `@Component`, `@RestController` and `@RequestMapping` without needing to use fully-qualified names or `import` statements.

Tip

Many Spring annotations will work without using `import` statements. Try running your application to see what fails before adding imports.

Automatic main method

Unlike the equivalent Java application, you do not need to include a `public static void main(String[] args)` method with your Groovy scripts. A `SpringApplication` is automatically created, with your compiled code acting as the `source`.

Custom dependency management

By default, the CLI uses the dependency management declared in `spring-boot-dependencies` when resolving `@Grab` dependencies. Additional dependency management, that will override the default dependency management, can be configured using the `@DependencyManagementBom` annotation. The annotation’s value should specify the coordinates (`groupId:artifactId:version`) of one or more Maven BOMs.

For example, the following declaration:

```
@DependencyManagementBom("com.example.custom-bom:1.0.0")
```

Will pick up `custom-bom-1.0.0.pom` in a Maven repository under `com/example/custom-versions/1.0.0/`.

When multiple BOMs are specified they are applied in the order that they’re declared. For example:

```
@DependencyManagementBom(["com.example.custom-bom:1.0.0",
    "com.example.another-bom:1.0.0"])
```

indicates that dependency management in `another-bom` will override the dependency management in `custom-bom`.

You can use `@DependencyManagementBom` anywhere that you can use `@Grab`, however, to ensure consistent ordering of the dependency management, you can only use `@DependencyManagementBom` at most once in your application. A useful source of dependency management (that is

a superset of Spring Boot’s dependency management) is the [Spring IO Platform](#), e.g. `@DependencyManagementBom('io.spring.platform:platform-bom:1.1.2.RELEASE')`.

63.2 Applications with multiple source files

You can use “shell globbing” with all commands that accept file input. This allows you to easily use multiple files from a single directory, e.g.

```
$ spring run *.groovy
```

63.3 Packaging your application

You can use the `jar` command to package your application into a self-contained executable jar file. For example:

```
$ spring jar my-app.jar *.groovy
```

The resulting jar will contain the classes produced by compiling the application and all of the application’s dependencies so that it can then be run using `java -jar`. The jar file will also contain entries from the application’s classpath. You can add explicit paths to the jar using `--include` and `--exclude` (both are comma-separated, and both accept prefixes to the values “+” and “-” to signify that they should be removed from the defaults). The default includes are

```
public/**, resources/**, static/**, templates/**, META-INF/**, *
```

and the default excludes are

```
.*, repository/**, build/**, target/**, **/*.jar, **/*.groovy
```

See the output of `spring help jar` for more information.

63.4 Initialize a new project

The `init` command allows you to create a new project using start.spring.io without leaving the shell. For example:

```
$ spring init --dependencies=web,data-jpa my-project
Using service at https://start.spring.io
Project extracted to '/Users/developer/example/my-project'
```

This creates a `my-project` directory with a Maven-based project using `spring-boot-starter-web` and `spring-boot-starter-data-jpa`. You can list the capabilities of the service using the `--list` flag

```
$ spring init --list
=====
Capabilities of https://start.spring.io
=====

Available dependencies:
-----
actuator - Actuator: Production ready features to help you monitor and manage your application
...
web - Web: Support for full-stack web development, including Tomcat and spring-webmvc
websocket - WebSocket: Support for WebSocket development
ws - WS: Support for Spring Web Services

Available project types:
```

```

-----
gradle-build - Gradle Config [format:build, build:gradle]
gradle-project - Gradle Project [format:project, build:gradle]
maven-build - Maven POM [format:build, build:maven]
maven-project - Maven Project [format:project, build:maven] (default)
...

```

The `init` command supports many options, check the `help` output for more details. For instance, the following command creates a gradle project using Java 8 and `war` packaging:

```

$ spring init --build=gradle --java-version=1.8 --dependencies=websocket --packaging=war sample-app.zip
Using service at https://start.spring.io
Content saved to 'sample-app.zip'

```

63.5 Using the embedded shell

Spring Boot includes command-line completion scripts for BASH and zsh shells. If you don't use either of these shells (perhaps you are a Windows user) then you can use the `shell` command to launch an integrated shell.

```

$ spring shell
Spring Boot (v2.0.0.M3)
Hit TAB to complete. Type \'help\' and hit RETURN for help, and \'exit\' to quit.

```

From inside the embedded shell you can run other commands directly:

```

$ version
Spring CLI v2.0.0.M3

```

The embedded shell supports ANSI color output as well as `tab` completion. If you need to run a native command you can use the `!` prefix. Hitting `ctrl-c` will exit the embedded shell.

63.6 Adding extensions to the CLI

You can add extensions to the CLI using the `install` command. The command takes one or more sets of artifact coordinates in the format `group:artifact:version`. For example:

```

$ spring install com.example:spring-boot-cli-extension:1.0.0.RELEASE

```

In addition to installing the artifacts identified by the coordinates you supply, all of the artifacts' dependencies will also be installed.

To uninstall a dependency use the `uninstall` command. As with the `install` command, it takes one or more sets of artifact coordinates in the format `group:artifact:version`. For example:

```

$ spring uninstall com.example:spring-boot-cli-extension:1.0.0.RELEASE

```

It will uninstall the artifacts identified by the coordinates you supply and their dependencies.

To uninstall all additional dependencies you can use the `--all` option. For example:

```

$ spring uninstall --all

```


64. Developing application with the Groovy beans DSL

Spring Framework 4.0 has native support for a `beans{}` “DSL” (borrowed from [Grails](#)), and you can embed bean definitions in your Groovy application scripts using the same format. This is sometimes a good way to include external features like middleware declarations. For example:

```
@Configuration
class Application implements CommandLineRunner {

    @Autowired
    SharedService service

    @Override
    void run(String... args) {
        println service.message
    }
}

import my.company.SharedService

beans {
    service(SharedService) {
        message = "Hello World"
    }
}
```

You can mix class declarations with `beans{}` in the same file as long as they stay at the top level, or you can put the beans DSL in a separate file if you prefer.

65. Configuring the CLI with settings.xml

The Spring Boot CLI uses Aether, Maven's dependency resolution engine, to resolve dependencies. The CLI makes use of the Maven configuration found in `~/.m2/settings.xml` to configure Aether. The following configuration settings are honored by the CLI:

- Offline
- Mirrors
- Servers
- Proxies
- Profiles
 - Activation
 - Repositories
- Active profiles

Please refer to [Maven's settings documentation](#) for further information.

66. What to read next

There are some [sample groovy scripts](#) available from the GitHub repository that you can use to try out the Spring Boot CLI. There is also extensive Javadoc throughout the [source code](#).

If you find that you reach the limit of the CLI tool, you will probably want to look at converting your application to full Gradle or Maven built “groovy project”. The next section covers Spring Boot’s [Build tool plugins](#) that you can use with Gradle or Maven.

Part VIII. Build tool plugins

Spring Boot provides build tool plugins for Maven and Gradle. The plugins offer a variety of features, including the packaging of executable jars. This section provides more details on both plugins, as well as some help should you need to extend an unsupported build system. If you are just getting started, you might want to read "[Chapter 13, *Build systems*](#)" from the [Part III, "Using Spring Boot"](#) section first.

67. Spring Boot Maven plugin

The [Spring Boot Maven Plugin](#) provides Spring Boot support in Maven, allowing you to package executable jar or war archives and run an application “in-place”. To use it you must be using Maven 3.2 (or better).

Note

Refer to the [Spring Boot Maven Plugin Site](#) for complete plugin documentation.

67.1 Including the plugin

To use the Spring Boot Maven Plugin simply include the appropriate XML in the `plugins` section of your `pom.xml`

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd">
  <modelVersion>4.0.0</modelVersion>
  <!-- ... -->
  <build>
    <plugins>
      <plugin>
        <groupId>org.springframework.boot</groupId>
        <artifactId>spring-boot-maven-plugin</artifactId>
        <version>2.0.0.M3</version>
        <executions>
          <execution>
            <goals>
              <goal>repackage</goal>
            </goals>
          </execution>
        </executions>
      </plugin>
    </plugins>
  </build>
</project>
```

This configuration will repackage a jar or war that is built during the `package` phase of the Maven lifecycle. The following example shows both the repackaged jar, as well as the original jar, in the `target` directory:

```
$ mvn package
$ ls target/*.jar
target/myproject-1.0.0.jar target/myproject-1.0.0.jar.original
```

If you don't include the `<execution/>` configuration as above, you can run the plugin on its own (but only if the `package` goal is used as well). For example:

```
$ mvn package spring-boot:repackage
$ ls target/*.jar
target/myproject-1.0.0.jar target/myproject-1.0.0.jar.original
```

If you are using a milestone or snapshot release you will also need to add appropriate `pluginRepository` elements:

```
<pluginRepositories>
  <pluginRepository>
    <id>spring-snapshots</id>
    <url>http://repo.spring.io/snapshot</url>
```

```

</pluginRepository>
<pluginRepository>
  <id>spring-milestones</id>
  <url>http://repo.spring.io/milestone</url>
</pluginRepository>
</pluginRepositories>

```

67.2 Packaging executable jar and war files

Once `spring-boot-maven-plugin` has been included in your `pom.xml` it will automatically attempt to rewrite archives to make them executable using the `spring-boot:repackage` goal. You should configure your project to build a jar or war (as appropriate) using the usual packaging element:

```

<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd">
  <!-- ... -->
  <packaging>jar</packaging>
  <!-- ... -->
</project>

```

Your existing archive will be enhanced by Spring Boot during the `package` phase. The main class that you want to launch can either be specified using a configuration option, or by adding a `Main-Class` attribute to the manifest in the usual way. If you don't specify a main class the plugin will search for a class with a public static void `main(String[] args)` method.

To build and run a project artifact, you can type the following:

```

$ mvn package
$ java -jar target/mymodule-0.0.1-SNAPSHOT.jar

```

To build a war file that is both executable and deployable into an external container you need to mark the embedded container dependencies as "provided", e.g:

```

<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd">
  <!-- ... -->
  <packaging>war</packaging>
  <!-- ... -->
  <dependencies>
    <dependency>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-starter-web</artifactId>
    </dependency>
    <dependency>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-starter-tomcat</artifactId>
      <scope>provided</scope>
    </dependency>
  </dependencies>
</project>

```

Tip

See the [“Section 86.1, “Create a deployable war file””](#) section for more details on how to create a deployable war file.

Advanced configuration options and examples are available in the [plugin info page](#).

68. Spring Boot Gradle plugin

The Spring Boot Gradle Plugin provides Spring Boot support in Gradle, allowing you to package executable jar or war archives, run Spring Boot applications and use the dependency management provided by `spring-boot-dependencies`. It requires Gradle 3.4 or later. Please refer to the plugin's documentation to learn more:

- Reference ([HTML](#) and [PDF](#))
- [API](#)

69. Spring Boot AntLib module

The Spring Boot AntLib module provides basic Spring Boot support for Apache Ant. You can use the module to create executable jars. To use the module you need to declare an additional `spring-boot` namespace in your `build.xml`:

```
<project xmlns:ivy="antlib:org.apache.ivy.ant"
  xmlns:spring-boot="antlib:org.springframework.boot.ant"
  name="myapp" default="build">
  ...
</project>
```

You'll need to remember to start Ant using the `-lib` option, for example:

```
$ ant -lib <folder containing spring-boot-antlib-2.0.0.M3.jar>
```

Tip

The “Using Spring Boot” section includes a more complete example of [using Apache Ant with `spring-boot-antlib`](#)

69.1 Spring Boot Ant tasks

Once the `spring-boot-antlib` namespace has been declared, the following additional tasks are available.

`spring-boot:exejar`

The `exejar` task can be used to create a Spring Boot executable jar. The following attributes are supported by the task:

Attribute	Description	Required
<code>destfile</code>	The destination jar file to create	Yes
<code>classes</code>	The root directory of Java class files	Yes
<code>start-class</code>	The main application class to run	No (<i>default is first class found declaring a main method</i>)

The following nested elements can be used with the task:

Element	Description
<code>resources</code>	One or more Resource Collections describing a set of Resources that should be added to the content of the created jar file.
<code>lib</code>	One or more Resource Collections that should be added to the set of jar libraries that make up the runtime dependency classpath of the application.

Examples

Specify `start-class`.


```
<spring-boot:exejar destfile="target/my-application.jar"
  classes="target/classes" start-class="com.foo.MyApplication">
  <resources>
    <fileset dir="src/main/resources" />
  </resources>
  <lib>
    <fileset dir="lib" />
  </lib>
</spring-boot:exejar>
```

Detect start-class.

```
<exejar destfile="target/my-application.jar" classes="target/classes">
  <lib>
    <fileset dir="lib" />
  </lib>
</exejar>
```

69.2 spring-boot:findmainclass

The `findmainclass` task is used internally by `exejar` to locate a class declaring a `main`. You can also use this task directly in your build if needed. The following attributes are supported

Attribute	Description	Required
<code>classesroot</code>	The root directory of Java class files	Yes (<i>unless mainclass is specified</i>)
<code>mainclass</code>	Can be used to short-circuit the <code>main</code> class search	No
<code>property</code>	The Ant property that should be set with the result	No (<i>result will be logged if unspecified</i>)

Examples

Find and log.

```
<findmainclass classesroot="target/classes" />
```

Find and set.

```
<findmainclass classesroot="target/classes" property="main-class" />
```

Override and set.

```
<findmainclass mainclass="com.foo.MainClass" property="main-class" />
```

70. Supporting other build systems

If you want to use a build tool other than Maven, Gradle or Ant, you will likely need to develop your own plugin. Executable jars need to follow a specific format and certain entries need to be written in an uncompressed form (see the [executable jar format](#) section in the appendix for details).

The Spring Boot Maven and Gradle plugins both make use of `spring-boot-loader-tools` to actually generate jars. You are also free to use this library directly yourself if you need to.

70.1 Repackaging archives

To repackage an existing archive so that it becomes a self-contained executable archive use `org.springframework.boot.loader.tools.Repackager`. The `Repackager` class takes a single constructor argument that refers to an existing jar or war archive. Use one of the two available `repackage()` methods to either replace the original file or write to a new destination. Various settings can also be configured on the `repackager` before it is run.

70.2 Nested libraries

When repackaging an archive you can include references to dependency files using the `org.springframework.boot.loader.tools.Libraries` interface. We don't provide any concrete implementations of `Libraries` here as they are usually build system specific.

If your archive already includes libraries you can use `Libraries.NONE`.

70.3 Finding a main class

If you don't use `Repackager.setMainClass()` to specify a main class, the repackager will use [ASM](#) to read class files and attempt to find a suitable class with a `public static void main(String[] args)` method. An exception is thrown if more than one candidate is found.

70.4 Example repackage implementation

Here is a typical example repackage:

```
Repackager repackager = new Repackager(sourceJarFile);
repackager.setBackupSource(false);
repackager.repackage(new Libraries() {
    @Override
    public void doWithLibraries(LibraryCallback callback) throws IOException {
        // Build system specific implementation, callback for each dependency
        // callback.library(new Library(nestedFile, LibraryScope.COMPILE));
    }
});
```

71. What to read next

If you're interested in how the build tool plugins work you can look at the [spring-boot-tools](#) module on GitHub. More technical details of the [executable jar format](#) are covered in the appendix.

If you have specific build-related questions you can check out the "[how-to](#)" guides.

Part IX. 'How-to' guides

This section provides answers to some common 'how do I do that...' type of questions that often arise when using Spring Boot. This is by no means an exhaustive list, but it does cover quite a lot.

If you are having a specific problem that we don't cover here, you might want to check out stackoverflow.com to see if someone has already provided an answer; this is also a great place to ask new questions (please use the `spring-boot` tag).

We're also more than happy to extend this section; If you want to add a 'how-to' you can send us a [pull request](#).

72. Spring Boot application

72.1 Create your own FailureAnalyzer

`FailureAnalyzer` is a great way to intercept an exception on startup and turn it into a human-readable message, wrapped into a `FailureAnalysis`. Spring Boot provides such analyzer for application context related exceptions, JSR-303 validations and more. It is actually very easy to create your own.

`AbstractFailureAnalyzer` is a convenient extension of `FailureAnalyzer` that checks the presence of a specified exception type in the exception to handle. You can extend from that so that your implementation gets a chance to handle the exception only when it is actually present. If for whatever reason you can't handle the exception, return `null` to give another implementation a chance to handle the exception.

`FailureAnalyzer` implementations are to be registered in a `META-INF/spring.factories`: the following registers `ProjectConstraintViolationFailureAnalyzer`:

```
org.springframework.boot.diagnostics.FailureAnalyzer=\
com.example.ProjectConstraintViolationFailureAnalyzer
```

72.2 Troubleshoot auto-configuration

The Spring Boot auto-configuration tries its best to 'do the right thing', but sometimes things fail and it can be hard to tell why.

There is a really useful `ConditionEvaluationReport` available in any Spring Boot `ApplicationContext`. You will see it if you enable `DEBUG` logging output. If you use the `spring-boot-actuator` there is also an `autoconfig` endpoint that renders the report in JSON. Use that to debug the application and see what features have been added (and which not) by Spring Boot at runtime.

Many more questions can be answered by looking at the source code and the Javadoc. Some rules of thumb:

- Look for classes called `*AutoConfiguration` and read their sources, in particular the `@Conditional*` annotations to find out what features they enable and when. Add `--debug` to the command line or a System property `-Ddebug` to get a log on the console of all the auto-configuration decisions that were made in your app. In a running Actuator app look at the `autoconfig` endpoint (`/autoconfig` or the JMX equivalent) for the same information.
- Look for classes that are `@ConfigurationProperties` (e.g. `ServerProperties`) and read from there the available external configuration options. The `@ConfigurationProperties` has a `name` attribute which acts as a prefix to external properties, thus `ServerProperties` has `prefix="server"` and its configuration properties are `server.port`, `server.address` etc. In a running Actuator app look at the `configprops` endpoint.
- Look for uses of the `bind` method on the `Binder` to pull configuration values explicitly out of the `Environment` in a relaxed manner. It is often is used with a prefix.
- Look for `@Value` annotations that bind directly to the `Environment`.
- Look for `@ConditionalOnExpression` annotations that switch features on and off in response to SpEL expressions, normally evaluated with placeholders resolved from the `Environment`.

72.3 Customize the Environment or ApplicationContext before it starts

A `SpringApplication` has `ApplicationListeners` and `ApplicationContextInitializers` that are used to apply customizations to the context or environment. Spring Boot loads a number of such customizations for use internally from `META-INF/spring.factories`. There is more than one way to register additional ones:

- Programmatically per application by calling the `addListeners` and `addInitializers` methods on `SpringApplication` before you run it.
- Declaratively per application by setting `context.initializer.classes` or `context.listener.classes`.
- Declaratively for all applications by adding a `META-INF/spring.factories` and packaging a jar file that the applications all use as a library.

The `SpringApplication` sends some special `ApplicationEvents` to the listeners (even some before the context is created), and then registers the listeners for events published by the `ApplicationContext` as well. See [Section 23.5, “Application events and listeners”](#) in the ‘Spring Boot features’ section for a complete list.

It is also possible to customize the `Environment` before the application context is refreshed using `EnvironmentPostProcessor`. Each implementation should be registered in `META-INF/spring.factories`:

```
org.springframework.boot.env.EnvironmentPostProcessor=com.example.YourEnvironmentPostProcessor
```

The implementation can load arbitrary files and add them to the `Environment`. For instance, this example loads a YAML configuration file from the classpath:

```
public class EnvironmentPostProcessorExample implements EnvironmentPostProcessor {

    private final YamlPropertySourceLoader loader = new YamlPropertySourceLoader();

    @Override
    public void postProcessEnvironment(ConfigurableEnvironment environment,
        SpringApplication application) {
        Resource path = new ClassPathResource("com/example/myapp/config.yml");
        PropertySource<?> propertySource = loadYaml(path);
        environment.getPropertySources().addLast(propertySource);
    }

    private PropertySource<?> loadYaml(Resource path) {
        if (!path.exists()) {
            throw new IllegalArgumentException("Resource " + path + " does not exist");
        }
        try {
            return this.loader.load("custom-resource", path, null);
        }
        catch (IOException ex) {
            throw new IllegalStateException(
                "Failed to load yaml configuration from " + path, ex);
        }
    }
}
```

Tip

The `Environment` will already have been prepared with all the usual property sources that Spring Boot loads by default. It is therefore possible to get the location of the file from the environment. This example adds the `custom-resource` property source at the end of the list so that a key defined in any of the usual other locations takes precedence. A custom implementation may obviously defines another order.

Note

While using `@PropertySource` on your `@SpringBootApplication` seems convenient and easy enough to load a custom resource in the `Environment`, we do not recommend it as Spring Boot prepares the `Environment` before the `ApplicationContext` is refreshed. Any key defined via `@PropertySource` will be loaded too late to have any effect on auto-configuration.

72.4 Build an `ApplicationContext` hierarchy (adding a parent or root context)

You can use the `ApplicationBuilder` class to create parent/child `ApplicationContext` hierarchies. See [Section 23.4, “Fluent builder API”](#) in the ‘Spring Boot features’ section for more information.

72.5 Create a non-web application

Not all Spring applications have to be web applications (or web services). If you want to execute some code in a `main` method, but also bootstrap a Spring application to set up the infrastructure to use, then it’s easy with the `SpringApplication` features of Spring Boot. A `SpringApplication` changes its `ApplicationContext` class depending on whether it thinks it needs a web application or not. The first thing you can do to help it is to just leave the servlet API dependencies off the classpath. If you can’t do that (e.g. you are running 2 applications from the same code base) then you can explicitly call `setWebEnvironment(false)` on your `SpringApplication` instance, or set the `applicationContextClass` property (through the Java API or with external properties). Application code that you want to run as your business logic can be implemented as a `CommandLineRunner` and dropped into the context as a `@Bean` definition.

73. Properties & configuration

73.1 Automatically expand properties at build time

Rather than hardcoding some properties that are also specified in your project's build configuration, you can automatically expand them using the existing build configuration instead. This is possible in both Maven and Gradle.

Automatic property expansion using Maven

You can automatically expand properties from the Maven project using resource filtering. If you use the `spring-boot-starter-parent` you can then refer to your Maven 'project properties' via `@. .@` placeholders, e.g.

```
app.encoding=@project.build.sourceEncoding@
app.java.version=@java.version@
```

Tip

The `spring-boot:run` can add `src/main/resources` directly to the classpath (for hot reloading purposes) if you enable the `addResources` flag. This circumvents the resource filtering and this feature. You can use the `exec:java` goal instead or customize the plugin's configuration, see the [plugin usage page](#) for more details.

If you don't use the starter parent, in your `pom.xml` you need (inside the `<build/>` element):

```
<resources>
  <resource>
    <directory>src/main/resources</directory>
    <filtering>true</filtering>
  </resource>
</resources>
```

and (inside `<plugins/>`):

```
<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-resources-plugin</artifactId>
  <version>2.7</version>
  <configuration>
    <delimiters>
      <delimiter>@</delimiter>
    </delimiters>
    <useDefaultDelimiters>>false</useDefaultDelimiters>
  </configuration>
</plugin>
```

Note

The `useDefaultDelimiters` property is important if you are using standard Spring placeholders in your configuration (e.g. `${foo}`). These may be expanded by the build if that property is not set to `false`.

Automatic property expansion using Gradle

You can automatically expand properties from the Gradle project by configuring the Java plugin's `processResources` task to do so:


```
processResources {
    expand(project.properties)
}
```

You can then refer to your Gradle project's properties via placeholders, e.g.

```
app.name=${name}
app.description=${description}
```

Note

Gradle's `expand` method uses Groovy's `SimpleTemplateEngine` which transforms `${..}` tokens. The `${..}` style conflicts with Spring's own property placeholder mechanism. To use Spring property placeholders together with automatic expansion the Spring property placeholders need to be escaped like `\${..}`.

73.2 Externalize the configuration of SpringApplication

A `SpringApplication` has bean properties (mainly setters) so you can use its Java API as you create the application to modify its behavior. Or you can externalize the configuration using properties in `spring.main.*`. E.g. in `application.properties` you might have.

```
spring.main.web-environment=false
spring.main.banner-mode=off
```

and then the Spring Boot banner will not be printed on startup, and the application will not be a web application.

Note

The example above also demonstrates how flexible binding allows the use of underscores (`_`) as well as dashes (`-`) in property names.

Properties defined in external configuration overrides the values specified via the Java API with the notable exception of the sources used to create the `ApplicationContext`. Let's consider this application

```
new SpringApplicationBuilder()
    .bannerMode(Banner.Mode.OFF)
    .sources(demo.MyApp.class)
    .run(args);
```

used with the following configuration:

```
spring.main.sources=com.acme.Config,com.acme.ExtraConfig
spring.main.banner-mode=console
```

The actual application will *now* show the banner (as overridden by configuration) and use three sources for the `ApplicationContext` (in that order): `demo.MyApp`, `com.acme.Config`, `com.acme.ExtraConfig`.

73.3 Change the location of external properties of an application

By default properties from different sources are added to the Spring `Environment` in a defined order (see [Chapter 24, Externalized Configuration](#) in the ‘Spring Boot features’ section for the exact order).

A nice way to augment and modify this is to add `@PropertySource` annotations to your application sources. Classes passed to the `SpringApplication` static convenience methods, and those added using `setSources()` are inspected to see if they have `@PropertySources`, and if they do, those properties are added to the `Environment` early enough to be used in all phases of the `ApplicationContext` lifecycle. Properties added in this way have lower priority than any added using the default locations (e.g. `application.properties`), system properties, environment variables or the command line.

You can also provide System properties (or environment variables) to change the behavior:

- `spring.config.name (SPRING_CONFIG_NAME)`, defaults to `application` as the root of the file name.
- `spring.config.location (SPRING_CONFIG_LOCATION)` is the file to load (e.g. a classpath resource or a URL). A separate `Environment` property source is set up for this document and it can be overridden by system properties, environment variables or the command line.

No matter what you set in the environment, Spring Boot will always load `application.properties` as described above. If YAML is used then files with the ‘.yml’ extension are also added to the list by default.

Spring Boot logs the configuration files that are loaded at `DEBUG` level and the candidates it has not found at `TRACE` level.

See [ConfigFileApplicationListener](#) for more detail.

73.4 Use ‘short’ command line arguments

Some people like to use (for example) `--port=9000` instead of `--server.port=9000` to set configuration properties on the command line. You can easily enable this by using placeholders in `application.properties`, e.g.

```
server.port=${port:8080}
```

Tip

If you are inheriting from the `spring-boot-starter-parent` POM, the default filter token of the `maven-resources-plugins` has been changed from `${*}` to `@` (i.e. `@maven.token@` instead of `${maven.token}`) to prevent conflicts with Spring-style placeholders. If you have enabled maven filtering for the `application.properties` directly, you may want to also change the default filter token to use [other delimiters](#).

Note

In this specific case the port binding will work in a PaaS environment like Heroku and Cloud Foundry, since in those two platforms the `PORT` environment variable is set automatically and Spring can bind to capitalized synonyms for `Environment` properties.

73.5 Use YAML for external properties

YAML is a superset of JSON and as such is a very convenient syntax for storing external properties in a hierarchical format. E.g.

```
spring:
  application:
    name: cruncher
  datasource:
    driverClassName: com.mysql.jdbc.Driver
    url: jdbc:mysql://localhost/test
  server:
    port: 9000
```

Create a file called `application.yml` and stick it in the root of your classpath, and also add `snakeyaml` to your dependencies (Maven coordinates `org.yaml:snakeyaml`, already included if you use the `spring-boot-starter`). A YAML file is parsed to a `Java Map<String, Object>` (like a JSON object), and Spring Boot flattens the map so that it is 1-level deep and has period-separated keys, a lot like people are used to with `Properties` files in Java.

The example YAML above corresponds to an `application.properties` file

```
spring.application.name=cruncher
spring.datasource.driverClassName=com.mysql.jdbc.Driver
spring.datasource.url=jdbc:mysql://localhost/test
server.port=9000
```

See [Section 24.6, “Using YAML instead of Properties”](#) in the ‘Spring Boot features’ section for more information about YAML.

73.6 Set the active Spring profiles

The Spring Environment has an API for this, but normally you would set a System property (`spring.profiles.active`) or an OS environment variable (`SPRING_PROFILES_ACTIVE`). E.g. launch your application with a `-D` argument (remember to put it before the main class or jar archive):

```
$ java -jar -Dspring.profiles.active=production demo-0.0.1-SNAPSHOT.jar
```

In Spring Boot you can also set the active profile in `application.properties`, e.g.

```
spring.profiles.active=production
```

A value set this way is replaced by the System property or environment variable setting, but not by the `SpringApplicationBuilder.profiles()` method. Thus the latter Java API can be used to augment the profiles without changing the defaults.

See [Chapter 25, Profiles](#) in the ‘Spring Boot features’ section for more information.

73.7 Change configuration depending on the environment

A YAML file is actually a sequence of documents separated by `---` lines, and each document is parsed separately to a flattened map.

If a YAML document contains a `spring.profiles` key, then the profiles value (comma-separated list of profiles) is fed into the `SpringEnvironment.acceptsProfiles()` and if any of those profiles is active that document is included in the final merge (otherwise not).

Example:

```
server:
  port: 9000
---
spring:
  profiles: development
server:
  port: 9001
---
spring:
  profiles: production
server:
  port: 0
```

In this example the default port is 9000, but if the Spring profile 'development' is active then the port is 9001, and if 'production' is active then it is 0.

The YAML documents are merged in the order they are encountered (so later values override earlier ones).

To do the same thing with properties files you can use `application-${profile}.properties` to specify profile-specific values.

73.8 Discover built-in options for external properties

Spring Boot binds external properties from `application.properties` (or `.yml`) (and other places) into an application at runtime. There is not (and technically cannot be) an exhaustive list of all supported properties in a single location because contributions can come from additional jar files on your classpath.

A running application with the Actuator features has a `configprops` endpoint that shows all the bound and bindable properties available through `@ConfigurationProperties`.

The appendix includes an [application.properties](#) example with a list of the most common properties supported by Spring Boot. The definitive list comes from searching the source code for `@ConfigurationProperties` and `@Value` annotations, as well as the occasional use of `Binder`.

74. Embedded Web servers

74.1 Use another Web server

The Spring Boot starters bring a default embedded container for you:

- `spring-boot-starter-web` brings Tomcat with `spring-boot-starter-tomcat`, but `spring-boot-starter-jetty` and `spring-boot-starter-undertow` can be used instead.
- `spring-boot-starter-webflux` brings Reactor Netty with `spring-boot-starter-reactor-netty`, but `spring-boot-starter-tomcat`, `spring-boot-starter-jetty` and `spring-boot-starter-undertow` can be used instead.

Note

Many starters only support Spring MVC, so they transitively bring `spring-boot-starter-web` into your application classpath

If you choose to use a different HTTP server, you need to exclude those dependencies and include the one you chose instead. Spring Boot provides separate starters for HTTP servers to help make this process as easy as possible.

Example in Maven, for Spring MVC:

```
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-web</artifactId>
  <exclusions>
    <!-- Exclude the Tomcat dependency -->
    <exclusion>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-starter-tomcat</artifactId>
    </exclusion>
  </exclusions>
</dependency>
<!-- Use Jetty instead -->
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-jetty</artifactId>
</dependency>
```

Example in Gradle, for Spring WebFlux:

```
configurations {
  // exclude Reactor Netty
  compile.exclude module: 'spring-boot-starter-reactor-netty'
}

dependencies {
  compile 'org.springframework.boot:spring-boot-starter-webflux'
  // Use Undertow instead
  compile 'org.springframework.boot:spring-boot-starter-undertow'
  // ...
}
```

Note

`spring-boot-starter-reactor-netty` is required to use the `webClient`, so excluding it is not required if you wish to use a different HTTP server.

74.2 Configure Jetty

Generally you can follow the advice from [Section 73.8, “Discover built-in options for external properties”](#) about `@ConfigurationProperties` (`ServerProperties` is the main one here), but also look at `ServletWebServerFactoryCustomizer`. The Jetty APIs are quite rich so once you have access to the `JettyServletWebServerFactory` you can modify it in a number of ways. Or the nuclear option is to add your own `JettyServletWebServerFactory`.

74.3 Add a Servlet, Filter or Listener to an application

There are two ways to add `Servlet`, `Filter`, `ServletContextListener` and the other listeners supported by the Servlet spec to your application. You can either provide Spring beans for them, or enable scanning for Servlet components.

Add a Servlet, Filter or Listener using a Spring bean

To add a `Servlet`, `Filter`, or `Servlet *Listener` provide a `@Bean` definition for it. This can be very useful when you want to inject configuration or dependencies. However, you must be very careful that they don't cause eager initialization of too many other beans because they have to be installed in the container very early in the application lifecycle (e.g. it's not a good idea to have them depend on your `DataSource` or JPA configuration). You can work around restrictions like that by initializing them lazily when first used instead of on initialization.

In the case of `Filters` and `Servlets` you can also add mappings and init parameters by adding a `FilterRegistrationBean` or `ServletRegistrationBean` instead of or as well as the underlying component.

Note

If no `dispatcherType` is specified on a filter registration, it will match `FORWARD`, `INCLUDE` and `REQUEST`. If `async` has been enabled, it will match `ASYNC` as well.

If you are migrating a filter that has no `dispatcher` element in `web.xml` you will need to specify a `dispatcherType` yourself:

```
@Bean
public FilterRegistrationBean myFilterRegistration() {
    FilterRegistrationBean registration = new FilterRegistrationBean();
    registration.setDispatcherTypes(DispatcherType.REQUEST);
    ....
    return registration;
}
```

Disable registration of a Servlet or Filter

As [described above](#) any `Servlet` or `Filter` beans will be registered with the servlet container automatically. To disable registration of a particular `Filter` or `Servlet` bean create a registration bean for it and mark it as disabled. For example:

```
@Bean
public FilterRegistrationBean registration(MyFilter filter) {
    FilterRegistrationBean registration = new FilterRegistrationBean(filter);
    registration.setEnabled(false);
    return registration;
}
```

Add Servlets, Filters, and Listeners using classpath scanning

`@WebServlet`, `@WebFilter`, and `@WebListener` annotated classes can be automatically registered with an embedded servlet container by annotating a `@Configuration` class with `@ServletComponentScan` and specifying the package(s) containing the components that you want to register. By default, `@ServletComponentScan` will scan from the package of the annotated class.

74.4 Change the HTTP port

In a standalone application the main HTTP port defaults to 8080, but can be set with `server.port` (e.g. in `application.properties` or as a System property). Thanks to relaxed binding of `Environment` values you can also use `SERVER_PORT` (e.g. as an OS environment variable).

To switch off the HTTP endpoints completely, but still create a `WebApplicationContext`, use `server.port=-1` (this is sometimes useful for testing).

For more details look at [the section called “Customizing embedded servlet containers”](#) in the ‘Spring Boot features’ section, or the [ServerProperties](#) source code.

74.5 Use a random unassigned HTTP port

To scan for a free port (using OS natives to prevent clashes) use `server.port=0`.

74.6 Discover the HTTP port at runtime

You can access the port the server is running on from log output or from the `ServletWebServerApplicationContext` via its `EmbeddedWebServer`. The best way to get that and be sure that it has initialized is to add a `@Bean` of type `ApplicationListener<ServletWebServerInitializedEvent>` and pull the container out of the event when it is published.

Tests that use `@SpringBootTest(webEnvironment=WebEnvironment.RANDOM_PORT)` can also inject the actual port into a field using the `@LocalServerPort` annotation. For example:

```
@RunWith(SpringJUnit4ClassRunner.class)
@SpringBootTest(webEnvironment=WebEnvironment.RANDOM_PORT)
public class MyWebIntegrationTests {

    @Autowired
    ServletWebServerApplicationContext server;

    @LocalServerPort
    int port;

    // ...

}
```

Note

`@LocalServerPort` is a meta-annotation for `@Value("${local.server.port}")`. Don't try to inject the port in a regular application. As we just saw, the value is only set once the container has initialized; contrary to a test, application code callbacks are processed early (i.e. before the value is actually available).

74.7 Configure SSL

SSL can be configured declaratively by setting the various `server.ssl.*` properties, typically in `application.properties` or `application.yml`. For example:

```
server.port=8443
server.ssl.key-store=classpath:keystore.jks
server.ssl.key-store-password=secret
server.ssl.key-password=another-secret
```

See [Ssl](#) for details of all of the supported properties.

Using configuration like the example above means the application will no longer support plain HTTP connector at port 8080. Spring Boot doesn't support the configuration of both an HTTP connector and an HTTPS connector via `application.properties`. If you want to have both then you'll need to configure one of them programmatically. It's recommended to use `application.properties` to configure HTTPS as the HTTP connector is the easier of the two to configure programmatically. See the [spring-boot-sample-tomcat-multi-connectors](#) sample project for an example.

74.8 Configure Access Logging

Access logs can be configured for Tomcat, Undertow and Jetty via their respective namespaces.

For instance, the following logs access on Tomcat with a [custom pattern](#).

```
server.tomcat.basedir=my-tomcat
server.tomcat.accesslog.enabled=true
server.tomcat.accesslog.pattern=%t %a "%r" %s (%D ms)
```

Note

The default location for logs is a `logs` directory relative to the tomcat base dir and said directory is a temp directory by default so you may want to fix Tomcat's base directory or use an absolute path for the logs. In the example above, the logs will be available in `my-tomcat/logs` relative to the working directory of the application.

Access logging for undertow can be configured in a similar fashion

```
server.undertow.accesslog.enabled=true
server.undertow.accesslog.pattern=%t %a "%r" %s (%D ms)
```

Logs are stored in a `logs` directory relative to the working directory of the application. This can be customized via `server.undertow.accesslog.directory`.

Finally, access logging for jetty can also be configured that way:

```
server.jetty.accesslog.enabled=true
server.jetty.accesslog.filename=/var/log/jetty-access.log
```

By default, logs will be redirected to `System.err`. For more details, please refer to [the documentation](#).

74.9 Use behind a front-end proxy server

Your application might need to send 302 redirects or render content with absolute links back to itself. When running behind a proxy, the caller wants a link to the proxy, and not to the physical address of

the machine hosting your app. Typically such situations are handled via a contract with the proxy, which will add headers to tell the back end how to construct links to itself.

If the proxy adds conventional `X-Forwarded-For` and `X-Forwarded-Proto` headers (most do this out of the box) the absolute links should be rendered correctly as long as `server.use-forward-headers` is set to `true` in your `application.properties`.

Note

If your application is running in Cloud Foundry or Heroku the `server.use-forward-headers` property will default to `true` if not specified. In all other instances it defaults to `false`.

Customize Tomcat's proxy configuration

If you are using Tomcat you can additionally configure the names of the headers used to carry “forwarded” information:

```
server.tomcat.remote-ip-header=x-your-remote-ip-header
server.tomcat.protocol-header=x-your-protocol-header
```

Tomcat is also configured with a default regular expression that matches internal proxies that are to be trusted. By default, IP addresses in `10/8`, `192.168/16`, `169.254/16` and `127/8` are trusted. You can customize the valve's configuration by adding an entry to `application.properties`, e.g.

```
server.tomcat.internal-proxies=192\\.168\\.\\d{1,3}\\.\\d{1,3}
```

Note

The double backslashes are only required when you're using a properties file for configuration. If you are using YAML, single backslashes are sufficient and a value that's equivalent to the one shown above would be `192\\.168\\.\\d{1,3}\\.\\d{1,3}`.

Note

You can trust all proxies by setting the `internal-proxies` to empty (but don't do this in production).

You can take complete control of the configuration of Tomcat's `RemoteIpValve` by switching the automatic one off (i.e. set `server.use-forward-headers=false`) and adding a new valve instance in a `TomcatServletWebServerFactory` bean.

74.10 Configure Tomcat

Generally you can follow the advice from [Section 73.8, “Discover built-in options for external properties”](#) about `@ConfigurationProperties` (`ServerProperties` is the main one here), but also look at `ServletWebServerFactoryCustomizer` and various Tomcat-specific `*Customizers` that you can add in one of those. The Tomcat APIs are quite rich so once you have access to the `TomcatServletWebServerFactory` you can modify it in a number of ways. Or the nuclear option is to add your own `TomcatServletWebServerFactory`.

74.11 Enable Multiple Connectors with Tomcat

Add a `org.apache.catalina.connector.Connector` to the `TomcatServletWebServerFactory` which can allow multiple connectors, e.g. HTTP and HTTPS connector:

```
@Bean
public ServletWebServerFactory servletContainer() {
    TomcatServletWebServerFactory tomcat = new TomcatServletWebServerFactory();
    tomcat.addAdditionalTomcatConnectors(createSslConnector());
    return tomcat;
}

private Connector createSslConnector() {
    Connector connector = new Connector("org.apache.coyote.http11.Http11NioProtocol");
    Http11NioProtocol protocol = (Http11NioProtocol) connector.getProtocolHandler();
    try {
        File keystore = new ClassPathResource("keystore").getFile();
        File truststore = new ClassPathResource("keystore").getFile();
        connector.setScheme("https");
        connector.setSecure(true);
        connector.setPort(8443);
        protocol.setSSLEnabled(true);
        protocol.setKeystoreFile(keystore.getAbsolutePath());
        protocol.setKeystorePass("changeit");
        protocol.setTruststoreFile(truststore.getAbsolutePath());
        protocol.setTruststorePass("changeit");
        protocol.setKeyAlias("apitester");
        return connector;
    }
    catch (IOException ex) {
        throw new IllegalStateException("can't access keystore: [" + "keystore"
            + "] or truststore: [" + "keystore" + "]", ex);
    }
}
```

74.12 Use Tomcat's LegacyCookieProcessor

The embedded Tomcat used by Spring Boot does not support "Version 0" of the Cookie format out of the box, and you may see the following error:

```
java.lang.IllegalArgumentException: An invalid character [32] was present in the Cookie value
```

If at all possible, you should consider updating your code to only store values compliant with later Cookie specifications. If, however, you're unable to change the way that cookies are written, you can instead configure Tomcat to use a `LegacyCookieProcessor`. To switch to the `LegacyCookieProcessor` use an `ServletWebServerFactoryCustomizer` bean that adds a `TomcatContextCustomizer`:

```
@Bean
public WebServerFactoryCustomizer<TomcatServletWebServerFactory> cookieProcessorCustomizer() {
    return (serverFactory) -> serverFactory.addContextCustomizers(
        (context) -> context.setCookieProcessor(new LegacyCookieProcessor());
    );
}
```

74.13 Configure Undertow

Generally you can follow the advice from [Section 73.8, "Discover built-in options for external properties"](#) about `@ConfigurationProperties` (`ServerProperties` and `ServerProperties.Undertow` are the main ones here), but also look at `ServletWebServerFactoryCustomizer`. Once you have access to the `UndertowServletWebServerFactory` you can use an

UndertowBuilderCustomizer to modify Undertow's configuration to meet your needs. Or the nuclear option is to add your own UndertowServletWebServerFactory.

74.14 Enable Multiple Listeners with Undertow

Add an UndertowBuilderCustomizer to the UndertowServletWebServerFactory and add a listener to the Builder:

```
@Bean
public UndertowServletWebServerFactory servletWebServerFactory() {
    UndertowServletWebServerFactory factory = new UndertowServletWebServerFactory();
    factory.addBuilderCustomizers(new UndertowBuilderCustomizer() {

        @Override
        public void customize(Builder builder) {
            builder.addHttpListener(8080, "0.0.0.0");
        }

    });
    return factory;
}
```

74.15 Create WebSocket endpoints using @ServerEndpoint

If you want to use @ServerEndpoint in a Spring Boot application that used an embedded container, you must declare a single ServerEndpointExporter @Bean:

```
@Bean
public ServerEndpointExporter serverEndpointExporter() {
    return new ServerEndpointExporter();
}
```

This bean will register any @ServerEndpoint annotated beans with the underlying WebSocket container. When deployed to a standalone servlet container this role is performed by a servlet container initializer and the ServerEndpointExporter bean is not required.

74.16 Enable HTTP response compression

HTTP response compression is supported by Jetty, Tomcat, and Undertow. It can be enabled via application.properties:

```
server.compression.enabled=true
```

By default, responses must be at least 2048 bytes in length for compression to be performed. This can be configured using the server.compression.min-response-size property.

By default, responses will only be compressed if their content type is one of the following:

- text/html
- text/xml
- text/plain
- text/css

This can be configured using the server.compression.mime-types property.

75. Spring MVC

75.1 Write a JSON REST service

Any Spring `@RestController` in a Spring Boot application should render JSON response by default as long as Jackson2 is on the classpath. For example:

```
@RestController
public class MyController {

    @RequestMapping("/thing")
    public MyThing thing() {
        return new MyThing();
    }
}
```

As long as `MyThing` can be serialized by Jackson2 (e.g. a normal POJO or Groovy object) then localhost:8080/thing will serve a JSON representation of it by default. Sometimes in a browser you might see XML responses because browsers tend to send accept headers that prefer XML.

75.2 Write an XML REST service

If you have the Jackson XML extension (`jackson-dataformat-xml`) on the classpath, it will be used to render XML responses and the very same example as we used for JSON would work. To use it, add the following dependency to your project:

```
<dependency>
  <groupId>com.fasterxml.jackson.dataformat</groupId>
  <artifactId>jackson-dataformat-xml</artifactId>
</dependency>
```

You may also want to add a dependency on Woodstox. It's faster than the default StAX implementation provided by the JDK and also adds pretty print support and improved namespace handling:

```
<dependency>
  <groupId>org.codehaus.woodstox</groupId>
  <artifactId>woodstox-core-asl</artifactId>
</dependency>
```

If Jackson's XML extension is not available, JAXB (provided by default in the JDK) will be used, with the additional requirement to have `MyThing` annotated as `@XmlElement`:

```
@XmlElement
public class MyThing {
    private String name;
    // .. getters and setters
}
```

To get the server to render XML instead of JSON you might have to send an `Accept: text/xml` header (or use a browser).

75.3 Customize the Jackson ObjectMapper

Spring MVC (client and server side) uses `HttpMessageConverters` to negotiate content conversion in an HTTP exchange. If Jackson is on the classpath you already get the default converter(s) provided by `Jackson2ObjectMapperBuilder`, an instance of which is auto-configured for you.

The `ObjectMapper` (or `XmlMapper` for Jackson XML converter) instance created by default has the following customized properties:

- `MapperFeature.DEFAULT_VIEW_INCLUSION` is disabled
- `DeserializationFeature.FAIL_ON_UNKNOWN_PROPERTIES` is disabled

Spring Boot has also some features to make it easier to customize this behavior.

You can configure the `ObjectMapper` and `XmlMapper` instances using the environment. Jackson provides an extensive suite of simple on/off features that can be used to configure various aspects of its processing. These features are described in six enums in Jackson which map onto properties in the environment:

Jackson enum	Environment property
<code>com.fasterxml.jackson.databind.DeserializationFeature</code>	<code>spring.jackson.deserialization.<feature_name>=true false</code>
<code>com.fasterxml.jackson.core.JsonGenerator.Feature</code>	<code>spring.jackson.generator.<feature_name>=true false</code>
<code>com.fasterxml.jackson.databind.MapperFeature</code>	<code>spring.jackson.mapper.<feature_name>=true false</code>
<code>com.fasterxml.jackson.core.JsonParser.Feature</code>	<code>spring.jackson.parser.<feature_name>=true false</code>
<code>com.fasterxml.jackson.databind.SerializationFeature</code>	<code>spring.jackson.serialization.<feature_name>=true false</code>
<code>com.fasterxml.jackson.annotation.JsonInclude.Include</code>	<code>spring.jackson.default-property-inclusion=always non_null non_absent non_default non_empty</code>

For example, to enable pretty print, set `spring.jackson.serialization.indent_output=true`. Note that, thanks to the use of [relaxed binding](#), the case of `indent_output` doesn't have to match the case of the corresponding enum constant which is `INDENT_OUTPUT`.

This environment-based configuration is applied to the auto-configured `Jackson2ObjectMapperBuilder` bean, and will apply to any mappers created using the builder, including the auto-configured `ObjectMapper` bean.

The context's `Jackson2ObjectMapperBuilder` can be customized by one or more `Jackson2ObjectMapperBuilderCustomizer` beans. Such customizer beans can be ordered and Boot's own customizer has an order of 0, allowing additional customization to be applied both before and after Boot's customization.

Any beans of type `com.fasterxml.jackson.databind.Module` will be automatically registered with the auto-configured `Jackson2ObjectMapperBuilder` and applied to any `ObjectMapper` instances that it creates. This provides a global mechanism for contributing custom modules when you add new features to your application.

If you want to replace the default `ObjectMapper` completely, either define a `@Bean` of that type and mark it as `@Primary`, or, if you prefer the builder-based approach, define a

`Jackson2ObjectMapperBuilder` `@Bean`. Note that in either case this will disable all auto-configuration of the `ObjectMapper`.

If you provide any `@Beans` of type `MappingJackson2HttpMessageConverter` then they will replace the default value in the MVC configuration. Also, a convenience bean is provided of type `HttpMessageConverters` (always available if you use the default MVC configuration) which has some useful methods to access the default and user-enhanced message converters.

See also the [Section 75.4, “Customize the `@ResponseBody` rendering”](#) section and the [WebMvcAutoConfiguration](#) source code for more details.

75.4 Customize the `@ResponseBody` rendering

Spring uses `HttpMessageConverters` to render `@ResponseBody` (or responses from `@RestController`). You can contribute additional converters by simply adding beans of that type in a Spring Boot context. If a bean you add is of a type that would have been included by default anyway (like `MappingJackson2HttpMessageConverter` for JSON conversions) then it will replace the default value. A convenience bean is provided of type `HttpMessageConverters` (always available if you use the default MVC configuration) which has some useful methods to access the default and user-enhanced message converters (useful, for example if you want to manually inject them into a custom `RestTemplate`).

As in normal MVC usage, any `WebMvcConfigurer` beans that you provide can also contribute converters by overriding the `configureMessageConverters` method, but unlike with normal MVC, you can supply only additional converters that you need (because Spring Boot uses the same mechanism to contribute its defaults). Finally, if you opt-out of the Spring Boot default MVC configuration by providing your own `@EnableWebMvc` configuration, then you can take control completely and do everything manually using `getMessageConverters` from `WebMvcConfigurationSupport`.

See the [WebMvcAutoConfiguration](#) source code for more details.

75.5 Handling Multipart File Uploads

Spring Boot embraces the Servlet 3 `javax.servlet.http.Part` API to support uploading files. By default Spring Boot configures Spring MVC with a maximum file of 1MB per file and a maximum of 10MB of file data in a single request. You may override these values, as well as the location to which intermediate data is stored (e.g., to the `/tmp` directory) and the threshold past which data is flushed to disk by using the properties exposed in the `MultipartProperties` class. If you want to specify that files be unlimited, for example, set the `spring.servlet.multipart.max-file-size` property to `-1`.

The multipart support is helpful when you want to receive multipart encoded file data as a `@RequestParam`-annotated parameter of type `MultipartFile` in a Spring MVC controller handler method.

See the [MultipartAutoConfiguration](#) source for more details.

75.6 Switch off the Spring MVC DispatcherServlet

Spring Boot wants to serve all content from the root of your application `/` down. If you would rather map your own servlet to that URL you can do it, but of course you may lose some of the other Boot MVC features. To add your own servlet and map it to the root resource just declare a `@Bean` of type `Servlet`

and give it the special bean name `dispatcherServlet` (You can also create a bean of a different type with that name if you want to switch it off and not replace it).

75.7 Switch off the Default MVC configuration

The easiest way to take complete control over MVC configuration is to provide your own `@Configuration` with the `@EnableWebMvc` annotation. This will leave all MVC configuration in your hands.

75.8 Customize ViewResolvers

A `ViewResolver` is a core component of Spring MVC, translating view names in `@Controller` to actual `View` implementations. Note that `ViewResolvers` are mainly used in UI applications, rather than REST-style services (a `View` is not used to render a `@ResponseBody`). There are many implementations of `ViewResolver` to choose from, and Spring on its own is not opinionated about which ones you should use. Spring Boot, on the other hand, installs one or two for you depending on what it finds on the classpath and in the application context. The `DispatcherServlet` uses all the resolvers it finds in the application context, trying each one in turn until it gets a result, so if you are adding your own you have to be aware of the order and in which position your resolver is added.

`WebMvcAutoConfiguration` adds the following `ViewResolvers` to your context:

- An `InternalResourceViewResolver` with bean id `'defaultViewResolver'`. This one locates physical resources that can be rendered using the `DefaultServlet` (e.g. static resources and JSP pages if you are using those). It applies a prefix and a suffix to the view name and then looks for a physical resource with that path in the servlet context (defaults are both empty, but accessible for external configuration via `spring.mvc.view.prefix` and `spring.mvc.view.suffix`). It can be overridden by providing a bean of the same type.
- A `BeanNameViewResolver` with id `'beanNameViewResolver'`. This is a useful member of the view resolver chain and will pick up any beans with the same name as the `View` being resolved. It shouldn't be necessary to override or replace it.
- A `ContentNegotiatingViewResolver` with id `'viewResolver'` is only added if there **are** actually beans of type `View` present. This is a 'master' resolver, delegating to all the others and attempting to find a match to the 'Accept' HTTP header sent by the client. There is a useful [blog about ContentNegotiatingViewResolver](#) that you might like to study to learn more, and also look at the source code for detail. You can switch off the auto-configured `ContentNegotiatingViewResolver` by defining a bean named `'viewResolver'`.
- If you use Thymeleaf you will also have a `ThymeleafViewResolver` with id `'thymeleafViewResolver'`. It looks for resources by surrounding the view name with a prefix and suffix (externalized to `spring.thymeleaf.prefix` and `spring.thymeleaf.suffix`, defaults `'classpath:/templates/'` and `'.html'` respectively). It can be overridden by providing a bean of the same name.
- If you use FreeMarker you will also have a `FreeMarkerViewResolver` with id `'freeMarkerViewResolver'`. It looks for resources in a loader path (externalized to `spring.freemarker.templateLoaderPath`, default `'classpath:/templates/'`) by surrounding the view name with a prefix and suffix (externalized to `spring.freemarker.prefix` and `spring.freemarker.suffix`, with empty and `'.ftl'` defaults respectively). It can be overridden by providing a bean of the same name.

- If you use Groovy templates (actually if `groovy-templates` is on your classpath) you will also have a `GroovyMarkupViewResolver` with id `'groovyMarkupViewResolver'`. It looks for resources in a loader path by surrounding the view name with a prefix and suffix (externalized to `spring.groovy.template.prefix` and `spring.groovy.template.suffix`, defaults `'classpath:/templates/'` and `'tpl'` respectively). It can be overridden by providing a bean of the same name.

Check out [WebMvcAutoConfiguration](#), [ThymeleafAutoConfiguration](#), [FreeMarkerAutoConfiguration](#) and [GroovyTemplateAutoConfiguration](#)

76. HTTP clients

76.1 Configure RestTemplate to use a proxy

As described in [Section 33.1, “RestTemplate customization”](#), a `RestTemplateCustomizer` can be used with `RestTemplateBuilder` to build a customized `RestTemplate`. This is the recommended approach for creating a `RestTemplate` configured to use a proxy.

The exact details of the proxy configuration depend on the underlying client request factory that is being used. Here’s an example of configuring `HttpComponentsClientRequestFactory` with an `HttpClient` that uses a proxy for all hosts except `192.168.0.5`.

```
static class ProxyCustomizer implements RestTemplateCustomizer {

    @Override
    public void customize(RestTemplate restTemplate) {
        HttpHost proxy = new HttpHost("proxy.example.com");
        HttpClient httpClient = HttpClientBuilder.create()
            .setRoutePlanner(new DefaultProxyRoutePlanner(proxy) {

                @Override
                public HttpHost determineProxy(HttpHost target,
                    HttpRequest request, HttpContext context)
                    throws HttpException {
                    if (target.getHostName().equals("192.168.0.5")) {
                        return null;
                    }
                    return super.determineProxy(target, request, context);
                }
            })
            .build();
        restTemplate.setRequestFactory(
            new HttpComponentsClientHttpRequestFactory(httpClient));
    }
}
```

77. Logging

Spring Boot has no mandatory logging dependency, except for the Commons Logging API, of which there are many implementations to choose from. To use [Logback](#) you need to include it and `jcl-over-slf4j` (which implements the Commons Logging API) on the classpath. The simplest way to do that is through the starters which all depend on `spring-boot-starter-logging`. For a web application you only need `spring-boot-starter-web` since it depends transitively on the logging starter. For example, using Maven:

```
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-web</artifactId>
</dependency>
```

Spring Boot has a `LoggingSystem` abstraction that attempts to configure logging based on the content of the classpath. If Logback is available it is the first choice.

If the only change you need to make to logging is to set the levels of various loggers then you can do that in `application.properties` using the "logging.level" prefix, e.g.

```
logging.level.org.springframework.web=DEBUG
logging.level.org.hibernate=ERROR
```

You can also set the location of a file to log to (in addition to the console) using "logging.file".

To configure the more fine-grained settings of a logging system you need to use the native configuration format supported by the `LoggingSystem` in question. By default Spring Boot picks up the native configuration from its default location for the system (e.g. `classpath:logback.xml` for Logback), but you can set the location of the config file using the "logging.config" property.

77.1 Configure Logback for logging

If you put a `logback.xml` in the root of your classpath it will be picked up from there (or `logback-spring.xml` to take advantage of the templating features provided by Boot). Spring Boot provides a default base configuration that you can include if you just want to set levels, e.g.

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
  <include resource="org/springframework/boot/logging/logback/base.xml"/>
  <logger name="org.springframework.web" level="DEBUG"/>
</configuration>
```

If you look at that `base.xml` in the `spring-boot` jar, you will see that it uses some useful `System` properties which the `LoggingSystem` takes care of creating for you. These are:

- `${PID}` the current process ID.
- `${LOG_FILE}` if `logging.file` was set in Boot's external configuration.
- `${LOG_PATH}` if `logging.path` was set (representing a directory for log files to live in).
- `${LOG_EXCEPTION_CONVERSION_WORD}` if `logging.exception-conversion-word` was set in Boot's external configuration.

Spring Boot also provides some nice ANSI colour terminal output on a console (but not in a log file) using a custom Logback converter. See the default `base.xml` configuration for details.

If Groovy is on the classpath you should be able to configure Logback with `logback.groovy` as well (it will be given preference if present).

Configure logback for file only output

If you want to disable console logging and write output only to a file you need a custom `logback-spring.xml` that imports `file-appender.xml` but not `console-appender.xml`:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
  <include resource="org/springframework/boot/logging/logback/defaults.xml" />
  <property name="LOG_FILE" value="${LOG_FILE:-${LOG_PATH:-${LOG_TEMP:-${java.io.tmpdir:-/tmp}}}}spring.log"/>
  <include resource="org/springframework/boot/logging/logback/file-appender.xml" />
  <root level="INFO">
    <appender-ref ref="FILE" />
  </root>
</configuration>
```

You also need to add `logging.file` to your `application.properties`:

```
logging.file=myapplication.log
```

77.2 Configure Log4j for logging

Spring Boot supports [Log4j 2](#) for logging configuration if it is on the classpath. If you are using the starters for assembling dependencies that means you have to exclude Logback and then include log4j 2 instead. If you aren't using the starters then you need to provide `jcl-over-slf4j` (at least) in addition to Log4j 2.

The simplest path is probably through the starters, even though it requires some jiggling with excludes, .e.g. in Maven:

```
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-web</artifactId>
</dependency>
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter</artifactId>
  <exclusions>
    <exclusion>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-starter-logging</artifactId>
    </exclusion>
  </exclusions>
</dependency>
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-log4j2</artifactId>
</dependency>
```

Note

The use of the Log4j starters gathers together the dependencies for common logging requirements (e.g. including having Tomcat use `java.util.logging` but configuring the output using Log4j 2). See the Actuator Log4j 2 samples for more detail and to see it in action.

Use YAML or JSON to configure Log4j 2

In addition to its default XML configuration format, Log4j 2 also supports YAML and JSON configuration files. To configure Log4j 2 to use an alternative configuration file format, add the appropriate dependencies to the classpath and name your configuration files to match your chosen file format:

Format	Dependencies	File names
YAML	<code>com.fasterxml.jackson.core:jackson-databind</code> <code>com.fasterxml.jackson.dataformat:jackson-dataformat-yaml</code>	<code>log4j2.yaml</code> <code>log4j2.yml</code>
JSON	<code>com.fasterxml.jackson.core:jackson-databind</code>	<code>log4j2.json</code> <code>log4j2.jsn</code>

78. Data Access

78.1 Configure a custom DataSource

To configure your own `DataSource` define a `@Bean` of that type in your configuration. Spring Boot will reuse your `DataSource` anywhere one is required, including database initialization. If you need to externalize some settings, you can easily bind your `DataSource` to the environment (see [the section called “Third-party configuration”](#)).

```
@Bean
@ConfigurationProperties(prefix="app.datasource")
public DataSource dataSource() {
    return new FancyDataSource();
}
```

```
app.datasource.url=jdbc:h2:mem:mydb
app.datasource.username=sa
app.datasource.pool-size=30
```

Assuming that your `FancyDataSource` has regular JavaBean properties for the url, the username and the pool size, these settings will be bound automatically before the `DataSource` is made available to other components. The regular [database initialization](#) will also happen (so the relevant sub-set of `spring.datasource.*` can still be used with your custom configuration).

You can apply the same principle if you are configuring a custom JNDI `DataSource`:

```
@Bean(destroyMethod=" ")
@ConfigurationProperties(prefix="app.datasource")
public DataSource dataSource() throws Exception {
    JndiDataSourceLookup dataSourceLookup = new JndiDataSourceLookup();
    return dataSourceLookup.getDataSource("java:comp/env/jdbc/YourDS");
}
```

Spring Boot also provides a utility builder class `DataSourceBuilder` that can be used to create one of the standard data sources (if it is on the classpath). The builder can detect the one to use based on what's available on the classpath. It also auto detects the driver based on the JDBC url.

```
@Bean
@ConfigurationProperties("app.datasource")
public DataSource dataSource() {
    return DataSourceBuilder.create().build();
}
```

To run an app with that `DataSource`, all that is needed really is the connection information; pool-specific settings can also be provided, check the implementation that is going to be used at runtime for more details.

```
app.datasource.url=jdbc:mysql://localhost/test
app.datasource.username=dbuser
app.datasource.password=dbpass
app.datasource.pool-size=30
```

There is a catch however. Because the actual type of the connection pool is not exposed, no keys are generated in the metadata for your custom `DataSource` and no completion is available in your IDE (The `DataSource` interface doesn't expose any property). Also, if you happen to have Hikari on the classpath, this basic setup will not work because Hikari has no `url` property (but a `jdbcUrl` property). You will have to rewrite your configuration as follows:

```
app.datasource.jdbc-url=jdbc:mysql://localhost/test
app.datasource.username=dbuser
app.datasource.password=dbpass
app.datasource.maximum-pool-size=30
```

You can fix that by forcing the connection pool to use and return a dedicated implementation rather than `DataSource`. You won't be able to change the implementation at runtime but the list of options will be explicit.

```
@Bean
@ConfigurationProperties("app.datasource")
public HikariDataSource dataSource() {
    return (HikariDataSource) DataSourceBuilder.create()
        .type(HikariDataSource.class).build();
}
```

You can even go further by leveraging what `DataSourceProperties` does for you, that is providing a default embedded database if no url is provided with a sensible username and password for it. You can easily initialize a `DataSourceBuilder` from the state of any `DataSourceProperties` so you could just as well inject the one Spring Boot creates automatically. However, that would split your configuration in two namespaces: url, username, password, type and driver on `spring.datasource` and the rest on your custom namespace (`app.datasource`). To avoid that, you can redefine a custom `DataSourceProperties` on your custom namespace:

```
@Bean
@Primary
@ConfigurationProperties("app.datasource")
public DataSourceProperties dataSourceProperties() {
    return new DataSourceProperties();
}

@Bean
@ConfigurationProperties("app.datasource")
public HikariDataSource dataSource(DataSourceProperties properties) {
    return (HikariDataSource) properties.initializeDataSourceBuilder()
        .type(HikariDataSource.class).build();
}
```

This setup puts you *in pair* with what Spring Boot does for you by default, except that a dedicated connection pool is chosen (in code) and its settings are exposed in the same namespace. Because `DataSourceProperties` is taking care of the `url/jdbcUrl` translation for you, you can configure it like this:

```
app.datasource.url=jdbc:mysql://localhost/test
app.datasource.username=dbuser
app.datasource.password=dbpass
app.datasource.maximum-pool-size=30
```

Note

Because your custom configuration chooses to go with Hikari, `app.datasource.type` will have no effect. In practice the builder will be initialized with whatever value you might set there and then overridden by the call to `.type()`.

See [Section 29.1, “Configure a DataSource”](#) in the ‘Spring Boot features’ section and the [DataSourceAutoConfiguration](#) class for more details.

78.2 Configure Two DataSources

If you need to configure multiple data sources, you can apply the same tricks that are described in the previous section. You must, however, mark one of the `DataSource @Primary` as various auto-configurations down the road expect to be able to get one by type.

If you create your own `DataSource`, the auto-configuration will back off. In the example below, we provide the *exact* same features set as what the auto-configuration provides on the primary data source:

```
@Bean
@Primary
@ConfigurationProperties("app.datasource.foo")
public DataSourceProperties fooDataSourceProperties() {
    return new DataSourceProperties();
}

@Bean
@Primary
@ConfigurationProperties("app.datasource.foo")
public DataSource fooDataSource() {
    return fooDataSourceProperties().initializeDataSourceBuilder().build();
}

@Bean
@ConfigurationProperties("app.datasource.bar")
public BasicDataSource barDataSource() {
    return (BasicDataSource) DataSourceBuilder.create()
        .type(BasicDataSource.class).build();
}
```

Tip

`fooDataSourceProperties` has to be flagged `@Primary` so that the database initializer feature uses your copy (should you use that).

Both data sources are also bound for advanced customizations. For instance you could configure them as follows:

```
app.datasource.foo.type=com.zaxxer.hikari.HikariDataSource
app.datasource.foo.maximum-pool-size=30

app.datasource.bar.url=jdbc:mysql://localhost/test
app.datasource.bar.username=dbuser
app.datasource.bar.password=dbpass
app.datasource.bar.max-total=30
```

Of course, you can apply the same concept to the secondary `DataSource` as well:

```
@Bean
@Primary
@ConfigurationProperties("app.datasource.foo")
public DataSourceProperties fooDataSourceProperties() {
    return new DataSourceProperties();
}

@Bean
@Primary
@ConfigurationProperties("app.datasource.foo")
public DataSource fooDataSource() {
    return fooDataSourceProperties().initializeDataSourceBuilder().build();
}

@Bean
```

```

@ConfigurationProperties("app.datasource.bar")
public DataSourceProperties barDataSourceProperties() {
    return new DataSourceProperties();
}

@Bean
@ConfigurationProperties("app.datasource.bar")
public DataSource barDataSource() {
    return barDataSourceProperties().initializeDataSourceBuilder().build();
}

```

This final example configures two data sources on custom namespaces with the same logic as what Spring Boot would do in auto-configuration.

78.3 Use Spring Data repositories

Spring Data can create implementations for you of `@Repository` interfaces of various flavors. Spring Boot will handle all of that for you as long as those `@Repositories` are included in the same package (or a sub-package) of your `@EnableAutoConfiguration` class.

For many applications all you will need is to put the right Spring Data dependencies on your classpath (there is a `spring-boot-starter-data-jpa` for JPA and a `spring-boot-starter-data-mongodb` for MongoDB), create some repository interfaces to handle your `@Entity` objects. Examples are in the [JPA sample](#) or the [MongoDB sample](#).

Spring Boot tries to guess the location of your `@Repository` definitions, based on the `@EnableAutoConfiguration` it finds. To get more control, use the `@EnableJpaRepositories` annotation (from Spring Data JPA).

78.4 Separate `@Entity` definitions from Spring configuration

Spring Boot tries to guess the location of your `@Entity` definitions, based on the `@EnableAutoConfiguration` it finds. To get more control, you can use the `@EntityScan` annotation, e.g.

```

@Configuration
@EnableAutoConfiguration
@EntityScan(basePackageClasses=City.class)
public class Application {

    //...

}

```

78.5 Configure JPA properties

Spring Data JPA already provides some vendor-independent configuration options (e.g. for SQL logging) and Spring Boot exposes those, and a few more for hibernate as external configuration properties. Some of them are automatically detected according to the context so you shouldn't have to set them.

The `spring.jpa.hibernate.ddl-auto` is a special case in that it has different defaults depending on whether you are using an embedded database (`create-drop`) or not (`none`). The dialect to use is also automatically detected based on the current `DataSource` but you can set `spring.jpa.database` yourself if you want to be explicit and bypass that check on startup.

Note

Specifying a database leads to the configuration of a well-defined Hibernate dialect. Several databases have more than one `Dialect` and this may not suit your need. In that case, you can either set `spring.jpa.database` to default to let Hibernate figure things out or set the dialect using the `spring.jpa.database-platform` property.

The most common options to set are:

```
spring.jpa.hibernate.naming.physical-strategy=com.example.MyPhysicalNamingStrategy
spring.jpa.show-sql=true
```

In addition all properties in `spring.jpa.properties.*` are passed through as normal JPA properties (with the prefix stripped) when the local `EntityManagerFactory` is created.

78.6 Configure Hibernate Naming Strategy

Hibernate uses [two different naming strategies](#) to map names from the object model to the corresponding database names. The fully qualified class name of the physical and implicit strategy implementations can be configured using the `spring.jpa.hibernate.naming.physical-strategy` and `spring.jpa.hibernate.naming.implicit-strategy` properties respectively.

Spring Boot configures the physical naming strategy with `SpringPhysicalNamingStrategy` by default. This implementation provides the same table structure as Hibernate 4: all dots are replaced by underscores and camel cases are replaced by underscores as well. By default, all table names are generated in lower case but it is possible to override that flag if your schema requires it.

Concretely, a `PhoneNumber` entity will be mapped to the `telephone_number` table.

If you'd rather use Hibernate 5's default instead, set the following property:

```
spring.jpa.hibernate.naming.physical-strategy=org.hibernate.boot.model.naming.PhysicalNamingStrategyStandardImpl
```

See [HibernateJpaAutoConfiguration](#) and [JpaBaseConfiguration](#) for more details.

78.7 Use a custom EntityManagerFactory

To take full control of the configuration of the `EntityManagerFactory`, you need to add a `@Bean` named 'entityManagerFactory'. Spring Boot auto-configuration switches off its entity manager based on the presence of a bean of that type.

78.8 Use Two EntityManagers

Even if the default `EntityManagerFactory` works fine, you will need to define a new one because otherwise the presence of the second bean of that type will switch off the default. To make it easy to do that you can use the convenient `EntityManagerBuilder` provided by Spring Boot, or if you prefer you can just use the `LocalContainerEntityManagerFactoryBean` directly from Spring ORM.

Example:

```
// add two data sources configured as above

@Bean
public LocalContainerEntityManagerFactoryBean customerEntityManagerFactory()
```

```

EntityManagerFactoryBuilder builder) {
    return builder
        .dataSource(customerDataSource())
        .packages(Customer.class)
        .persistenceUnit("customers")
        .build();
}

@Bean
public LocalContainerEntityManagerFactoryBean orderEntityManagerFactory(
    EntityManagerFactoryBuilder builder) {
    return builder
        .dataSource(orderDataSource())
        .packages(Order.class)
        .persistenceUnit("orders")
        .build();
}

```

The configuration above almost works on its own. To complete the picture you need to configure `TransactionManagers` for the two `EntityManager`s as well. One of them could be picked up by the default `JpaTransactionManager` in Spring Boot if you mark it as `@Primary`. The other would have to be explicitly injected into a new instance. Or you might be able to use a JTA transaction manager spanning both.

If you are using Spring Data, you need to configure `@EnableJpaRepositories` accordingly:

```

@Configuration
@EnableJpaRepositories(basePackageClasses = Customer.class,
    entityManagerFactoryRef = "customerEntityManagerFactory")
public class CustomerConfiguration {
    ...
}

@Configuration
@EnableJpaRepositories(basePackageClasses = Order.class,
    entityManagerFactoryRef = "orderEntityManagerFactory")
public class OrderConfiguration {
    ...
}

```

78.9 Use a traditional persistence.xml

Spring doesn't require the use of XML to configure the JPA provider, and Spring Boot assumes you want to take advantage of that feature. If you prefer to use `persistence.xml` then you need to define your own `@Bean` of type `LocalEntityManagerFactoryBean` (with id `'entityManagerFactory'`, and set the persistence unit name there.

See [JpaBaseConfiguration](#) for the default settings.

78.10 Use Spring Data JPA and Mongo repositories

Spring Data JPA and Spring Data Mongo can both create `Repository` implementations for you automatically. If they are both present on the classpath, you might have to do some extra configuration to tell Spring Boot which one (or both) you want to create repositories for you. The most explicit way to do that is to use the standard Spring Data `@Enable*Repositories` and tell it the location of your `Repository` interfaces (where `'*'` is `'Jpa'` or `'Mongo'` or both).

There are also flags `spring.data.*.repositories.enabled` that you can use to switch the auto-configured repositories on and off in external configuration. This is useful for instance in case you want to switch off the Mongo repositories and still use the auto-configured `MongoTemplate`.

The same obstacle and the same features exist for other auto-configured Spring Data repository types (Elasticsearch, Solr). Just change the names of the annotations and flags respectively.

78.11 Expose Spring Data repositories as REST endpoint

Spring Data REST can expose the `Repository` implementations as REST endpoints for you as long as Spring MVC has been enabled for the application.

Spring Boot exposes a set of useful properties from the `spring.data.rest` namespace that customize the `RepositoryRestConfiguration`. If you need to provide additional customization, you should use a `RepositoryRestConfigurer` bean.

Note

If you don't specify any order on your custom `RepositoryRestConfigurer` it will run after the one Spring Boot uses internally. If you need to specify an order, make sure it is higher than 0.

78.12 Configure a component that is used by JPA

If you want to configure a component that will be used by JPA then you need to ensure that the component is initialized before JPA. Where the component is auto-configured Spring Boot will take care of this for you. For example, when Flyway is auto-configured, Hibernate is configured to depend upon Flyway so that the latter has a chance to initialize the database before Hibernate tries to use it.

If you are configuring a component yourself, you can use an `EntityManagerFactoryDependsOnPostProcessor` subclass as a convenient way of setting up the necessary dependencies. For example, if you are using Hibernate Search with Elasticsearch as its index manager then any `EntityManagerFactory` beans must be configured to depend on the `elasticsearchClient` bean:

```
/**
 * {@link EntityManagerFactoryDependsOnPostProcessor} that ensures that
 * {@link EntityManagerFactory} beans depend on the elasticsearchClient bean.
 */
@Configuration
static class ElasticsearchJpaDependencyConfiguration
    extends EntityManagerFactoryDependsOnPostProcessor {

    ElasticsearchJpaDependencyConfiguration() {
        super("elasticsearchClient");
    }
}
```

79. Database initialization

An SQL database can be initialized in different ways depending on what your stack is. Or of course you can do it manually as long as the database is a separate process.

79.1 Initialize a database using JPA

JPA has features for DDL generation, and these can be set up to run on startup against the database. This is controlled through two external properties:

- `spring.jpa.generate-ddl` (boolean) switches the feature on and off and is vendor independent.
- `spring.jpa.hibernate.ddl-auto` (enum) is a Hibernate feature that controls the behavior in a more fine-grained way. See below for more detail.

79.2 Initialize a database using Hibernate

You can set `spring.jpa.hibernate.ddl-auto` explicitly and the standard Hibernate property values are `none`, `validate`, `update`, `create`, `create-drop`. Spring Boot chooses a default value for you based on whether it thinks your database is embedded (default `create-drop`) or not (default `none`). An embedded database is detected by looking at the `Connection` type: `hsqldb`, `h2` and `derby` are embedded, the rest are not. Be careful when switching from in-memory to a 'real' database that you don't make assumptions about the existence of the tables and data in the new platform. You either have to set `ddl-auto` explicitly, or use one of the other mechanisms to initialize the database.

Note

You can output the schema creation by enabling the `org.hibernate.SQL` logger. This is done for you automatically if you enable the [debug mode](#).

In addition, a file named `import.sql` in the root of the classpath will be executed on startup if Hibernate creates the schema from scratch (that is if the `ddl-auto` property is set to `create` or `create-drop`). This can be useful for demos and for testing if you are careful, but probably not something you want to be on the classpath in production. It is a Hibernate feature (nothing to do with Spring).

79.3 Initialize a database using Spring JDBC

Spring JDBC has a `DataSource` initializer feature. Spring Boot enables it by default and loads SQL from the standard locations `schema.sql` and `data.sql` (in the root of the classpath). In addition Spring Boot will load the `schema-${platform}.sql` and `data-${platform}.sql` files (if present), where `platform` is the value of `spring.datasource.platform`, e.g. you might choose to set it to the vendor name of the database (`hsqldb`, `h2`, `oracle`, `mysql`, `postgres` etc.). Spring Boot enables the fail-fast feature of the Spring JDBC initializer by default, so if the scripts cause exceptions the application will fail to start. The script locations can be changed by setting `spring.datasource.schema` and `spring.datasource.data`, and neither location will be processed if `spring.datasource.initialize=false`.

To disable the fail-fast you can set `spring.datasource.continue-on-error=true`. This can be useful once an application has matured and been deployed a few times, since the scripts can act as 'poor man's migrations' — inserts that fail mean that the data is already there, so there would be no need to prevent the application from running, for instance.

If you want to use the `schema.sql` initialization in a JPA app (with Hibernate) then `ddl-auto=create-drop` will lead to errors if Hibernate tries to create the same tables. To avoid those errors set `ddl-auto` explicitly to `""` (preferable) or `"none"`. Whether or not you use `ddl-auto=create-drop` you can always use `data.sql` to initialize new data.

79.4 Initialize a Spring Batch database

If you are using Spring Batch then it comes pre-packaged with SQL initialization scripts for most popular database platforms. Spring Boot will detect your database type, and execute those scripts by default, and in this case will switch the fail fast setting to false (errors are logged but do not prevent the application from starting). This is because the scripts are known to be reliable and generally do not contain bugs, so errors are ignorable, and ignoring them makes the scripts idempotent. You can switch off the initialization explicitly using `spring.batch.initializer.enabled=false`.

79.5 Use a higher-level database migration tool

Spring Boot supports two higher-level migration tools: [Flyway](#) and [Liquibase](#).

Execute Flyway database migrations on startup

To automatically run Flyway database migrations on startup, add the `org.flywaydb:flyway-core` to your classpath.

The migrations are scripts in the form `V<VERSION>__<NAME>.sql` (with `<VERSION>` an underscore-separated version, e.g. '1' or '2_1'). By default they live in a folder `classpath:db/migration` but you can modify that using `flyway.locations`. You can also add a special `{vendor}` placeholder to use vendor-specific scripts. Assume the following:

```
flyway.locations=db/migration/{vendor}
```

Rather than using `db/migration`, this configuration will set the folder to use according to the type of the database (i.e. `db/migration/mysql` for MySQL). The list of supported database are available in [DatabaseDriver](#).

See also the Flyway class from `flyway-core` for details of available settings like schemas etc. In addition Spring Boot provides a small set of properties in [FlywayProperties](#) that can be used to disable the migrations, or switch off the location checking. Spring Boot will call `Flyway.migrate()` to perform the database migration. If you would like more control, provide a `@Bean` that implements [FlywayMigrationStrategy](#).

Flyway supports SQL and Java [callbacks](#). To use SQL-based callbacks, place the callback scripts in the `classpath:db/migration` folder. To use Java-based callbacks, create one or more beans that implement `FlywayCallback` or, preferably, extend `BaseFlywayCallback`. Any such beans will be automatically registered with Flyway. They can be ordered using `@Order` or by implementing `Ordered`.

By default Flyway will autowire the (`@Primary`) `DataSource` in your context and use that for migrations. If you like to use a different `DataSource` you can create one and mark its `@Bean` as `@FlywayDataSource` - if you do that remember to create another one and mark it as `@Primary` if you want two data sources. Or you can use Flyway's native `DataSource` by setting `flyway.[url,user,password]` in external properties.

There is a [Flyway sample](#) so you can see how to set things up.

You can also use Flyway to provide data for specific scenarios. For example, you can place test-specific migrations in `src/test/resources` and they will only be run when your application starts for testing. If you want to be more sophisticated you can use profile-specific configuration to customize `flyway.locations` so that certain migrations will only run when a particular profile is active. For example, in `application-dev.properties` you could set `flyway.locations` to `classpath:/db/migration, classpath:/dev/db/migration` and migrations in `dev/db/migration` will only run when the `dev` profile is active.

Execute Liquibase database migrations on startup

To automatically run Liquibase database migrations on startup, add the `org.liquibase:liquibase-core` to your classpath.

The master change log is by default read from `db/changelog/db.changelog-master.yaml` but can be set using `liquibase.change-log`. In addition to YAML, Liquibase also supports JSON, XML, and SQL change log formats.

By default Liquibase will autowire the (`@Primary`) `DataSource` in your context and use that for migrations. If you like to use a different `DataSource` you can create one and mark its `@Bean` as `@LiquibaseDataSource` - if you do that remember to create another one and mark it as `@Primary` if you want two data sources. Or you can use Liquibase's native `DataSource` by setting `liquibase.[url,user,password]` in external properties.

See [LiquibaseProperties](#) for details of available settings like contexts, default schema etc.

There is a [Liquibase sample](#) so you can see how to set things up.

80. Messaging

80.1 Disable transacted JMS session

If your JMS broker does not support transacted session, you will have to disable the support of transactions altogether. If you create your own `JmsListenerContainerFactory` there is nothing to do since it won't be transacted by default. If you want to use the `DefaultJmsListenerContainerFactoryConfigurer` to reuse Spring Boot's default, you can disable transacted session as follows:

```
@Bean
public DefaultJmsListenerContainerFactory jmsListenerContainerFactory(
    ConnectionFactory connectionFactory,
    DefaultJmsListenerContainerFactoryConfigurer configurer) {
    DefaultJmsListenerContainerFactory listenerFactory =
        new DefaultJmsListenerContainerFactory();
    configurer.configure(listenerFactory, connectionFactory);
    listenerFactory.setTransactionManager(null);
    listenerFactory.setSessionTransacted(false);
    return listenerFactory;
}
```

This overrides the default factory and this should be applied to any other factory that your application defines, if any.

81. Batch applications

Note

By default, batch applications require a `DataSource` to store job details. If you want to deviate from that, you'll need to implement `BatchConfigurer`, see [The Javadoc of `@EnableBatchProcessing`](#) for more details.

81.1 Execute Spring Batch jobs on startup

Spring Batch auto-configuration is enabled by adding `@EnableBatchProcessing` (from Spring Batch) somewhere in your context.

By default it executes **all** `Jobs` in the application context on startup (see [JobLauncherCommandLineRunner](#) for details). You can narrow down to a specific job or jobs by specifying `spring.batch.job.names` (comma-separated job name patterns).

If the application context includes a `JobRegistry` then the jobs in `spring.batch.job.names` are looked up in the registry instead of being autowired from the context. This is a common pattern with more complex systems where multiple jobs are defined in child contexts and registered centrally.

See [BatchAutoConfiguration](#) and [@EnableBatchProcessing](#) for more details.

82. Actuator

82.1 Change the HTTP port or address of the actuator endpoints

In a standalone application the Actuator HTTP port defaults to the same as the main HTTP port. To make the application listen on a different port set the external property `management.port`. To listen on a completely different network address (e.g. if you have an internal network for management and an external one for user applications) you can also set `management.address` to a valid IP address that the server is able to bind to.

For more detail look at the [ManagementServerProperties](#) source code and [Section 50.3, “Customizing the management server port”](#) in the ‘Production-ready features’ section.

82.2 Customize the ‘whitelabel’ error page

Spring Boot installs a ‘whitelabel’ error page that you will see in browser client if you encounter a server error (machine clients consuming JSON and other media types should see a sensible response with the right error code).

Note

Set `server.error.whitelabel.enabled=false` to switch the default error page off which will restore the default of the servlet container that you are using. Note that Spring Boot will still attempt to resolve the error view so you’d probably add your own error page rather than disabling it completely.

Overriding the error page with your own depends on the templating technology that you are using. For example, if you are using Thymeleaf you would add an `error.html` template and if you are using FreeMarker you would add an `error.ftl` template. In general what you need is a View that resolves with a name of `error`, and/or a @Controller that handles the `/error` path. Unless you replaced some of the default configuration you should find a `BeanNameViewResolver` in your `ApplicationContext` so a @Bean with id `error` would be a simple way of doing that. Look at [ErrorMvcAutoConfiguration](#) for more options.

See also the section on [Error Handling](#) for details of how to register handlers in the servlet container.

82.3 Actuator and Jersey

Actuator HTTP endpoints are only available for Spring MVC-based applications. If you want to use Jersey and still use the actuator you will need to enable Spring MVC (by depending on `spring-boot-starter-web`, for example). By default, both Jersey and the Spring MVC dispatcher servlet are mapped to the same path (`/`). You will need to change the path for one of them (by configuring `server.servlet.path` for Spring MVC or `spring.jersey.application-path` for Jersey). For example, if you add `server.servlet.path=/system` into `application.properties`, the actuator HTTP endpoints will be available under `/system`.

83. Security

83.1 Switch off the Spring Boot security configuration

If you define a `@Configuration` with `@EnableWebSecurity` anywhere in your application it will switch off the default webapp security settings in Spring Boot (but leave the Actuator's security enabled). To tweak the defaults try setting properties in `security.*` (see [SecurityProperties](#) for details of available settings) and `SECURITY` section of [Common application properties](#).

83.2 Change the AuthenticationManager and add user accounts

If you provide a `@Bean` of type `AuthenticationManager` the default one will not be created, so you have the full feature set of Spring Security available (e.g. [various authentication options](#)).

Spring Security also provides a convenient `AuthenticationManagerBuilder` which can be used to build an `AuthenticationManager` with common options. The recommended way to use this in a webapp is to inject it into a void method in a `WebSecurityConfigurerAdapter`, e.g.

```
@Configuration
public class SecurityConfiguration extends WebSecurityConfigurerAdapter {

    @Autowired
    public void configureGlobal(AuthenticationManagerBuilder auth) throws Exception {
        auth.inMemoryAuthentication()
            .withUser("barry").password("password").roles("USER"); // ... etc.
    }

    // ... other stuff for application security
}
```

You will get the best results if you put this in a nested class, or a standalone class (i.e. not mixed in with a lot of other `@Beans` that might be allowed to influence the order of instantiation). The [secure web sample](#) is a useful template to follow.

If you experience instantiation issues (e.g. using JDBC or JPA for the user detail store) it might be worth extracting the `AuthenticationManagerBuilder` callback into a `GlobalAuthenticationConfigurerAdapter` (in the `init()` method so it happens before the authentication manager is needed elsewhere), e.g.

```
@Configuration
public class AuthenticationManagerConfiguration extends
    GlobalAuthenticationConfigurerAdapter {

    @Override
    public void init(AuthenticationManagerBuilder auth) {
        auth.inMemoryAuthentication() // ... etc.
    }

}
```

83.3 Enable HTTPS when running behind a proxy server

Ensuring that all your main endpoints are only available over HTTPS is an important chore for any application. If you are using Tomcat as a servlet container, then Spring Boot will add Tomcat's own

`RemoteIpValve` automatically if it detects some environment settings, and you should be able to rely on the `HttpServletRequest` to report whether it is secure or not (even downstream of a proxy server that handles the real SSL termination). The standard behavior is determined by the presence or absence of certain request headers (`x-forwarded-for` and `x-forwarded-proto`), whose names are conventional, so it should work with most front end proxies. You can switch on the valve by adding some entries to `application.properties`, e.g.

```
server.tomcat.remote_ip_header=x-forwarded-for
server.tomcat.protocol_header=x-forwarded-proto
```

(The presence of either of those properties will switch on the valve. Or you can add the `RemoteIpValve` yourself by adding a `TomcatServletWebServerFactory` bean.)

Spring Security can also be configured to require a secure channel for all (or some requests). To switch that on in a Spring Boot application you just need to set `security.require_ssl` to `true` in `application.properties`.

84. Hot swapping

84.1 Reload static content

There are several options for hot reloading. The recommended approach is to use [spring-boot-devtools](#) as it provides additional development-time features such as support for fast application restarts and LiveReload as well as sensible development-time configuration (e.g. template caching). Devtools works by monitoring the classpath for changes. This means that static resource changes must be "built" for the change to take affect. By default, this happens automatically in Eclipse when you save your changes. In IntelliJ IDEA, Make Project will trigger the necessary build. Due to the [default restart exclusions](#), changes to static resources will not trigger a restart of your application. They will, however, trigger a live reload.

Alternatively, running in an IDE (especially with debugging on) is a good way to do development (all modern IDEs allow reloading of static resources and usually also hot-swapping of Java class changes).

Finally, the [Maven and Gradle plugins](#) can be configured (see the `addResources` property) to support running from the command line with reloading of static files directly from source. You can use that with an external css/js compiler process if you are writing that code with higher level tools.

84.2 Reload templates without restarting the container

Most of the templating technologies supported by Spring Boot include a configuration option to disable caching (see below for details). If you're using the `spring-boot-devtools` module these properties will be [automatically configured](#) for you at development time.

Thymeleaf templates

If you are using Thymeleaf, then set `spring.thymeleaf.cache` to `false`. See [ThymeleafAutoConfiguration](#) for other Thymeleaf customization options.

FreeMarker templates

If you are using FreeMarker, then set `spring.freemarker.cache` to `false`. See [FreeMarkerAutoConfiguration](#) for other FreeMarker customization options.

Groovy templates

If you are using Groovy templates, then set `spring.groovy.template.cache` to `false`. See [GroovyTemplateAutoConfiguration](#) for other Groovy customization options.

84.3 Fast application restarts

The `spring-boot-devtools` module includes support for automatic application restarts. Whilst not as fast as technologies such as [JRebel](#) it's usually significantly faster than a "cold start". You should probably give it a try before investigating some of the more complex reload options discussed below.

For more details see the [Chapter 20, Developer tools](#) section.

84.4 Reload Java classes without restarting the container

Modern IDEs (Eclipse, IDEA, etc.) all support hot swapping of bytecode, so if you make a change that doesn't affect class or method signatures it should reload cleanly with no side effects.

85. Build

85.1 Generate build information

Both the Maven and Gradle plugin allow to generate build information containing the coordinates, name and version of the project. The plugin can also be configured to add additional properties through configuration. When such file is present, Spring Boot auto-configures a `BuildProperties` bean.

To generate build information with Maven, add an execution for the `build-info` goal:

```
<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
      <version>2.0.0.M3</version>
      <executions>
        <execution>
          <goals>
            <goal>build-info</goal>
          </goals>
        </execution>
      </executions>
    </plugin>
  </plugins>
</build>
```

Tip

Check the [Spring Boot Maven Plugin documentation](#) for more details.

And to do the same with Gradle:

```
springBoot {
    buildInfo()
}
```

Additional properties can be added using the DSL:

```
springBoot {
    buildInfo {
        additionalProperties = [
            'foo': 'bar'
        ]
    }
}
```

85.2 Generate git information

Both Maven and Gradle allow to generate a `git.properties` file containing information about the state of your `git` source code repository when the project was built.

For Maven users the `spring-boot-starter-parent` POM includes a pre-configured plugin to generate a `git.properties` file. Simply add the following declaration to your POM:

```
<build>
  <plugins>
    <plugin>
      <groupId>pl.project13.maven</groupId>
```

```

    <artifactId>git-commit-id-plugin</artifactId>
  </plugin>
</plugins>
</build>

```

Gradle users can achieve the same result using the [gradle-git-properties](#) plugin

```

plugins {
    id "com.gorylenko.gradle-git-properties" version "1.4.17"
}

```

85.3 Customize dependency versions

If you use a Maven build that inherits directly or indirectly from `spring-boot-dependencies` (for instance `spring-boot-starter-parent`) but you want to override a specific third-party dependency you can add appropriate `<properties>` elements. Browse the [spring-boot-dependencies](#) POM for a complete list of properties. For example, to pick a different `slf4j` version you would add the following:

```

<properties>
  <slf4j.version>1.7.5</slf4j.version>
</properties>

```

Note

This only works if your Maven project inherits (directly or indirectly) from `spring-boot-dependencies`. If you have added `spring-boot-dependencies` in your own `dependencyManagement` section with `<scope>import</scope>` you have to redefine the artifact yourself instead of overriding the property.

Warning

Each Spring Boot release is designed and tested against a specific set of third-party dependencies. Overriding versions may cause compatibility issues.

85.4 Create an executable JAR with Maven

The `spring-boot-maven-plugin` can be used to create an executable ‘fat’ JAR. If you are using the `spring-boot-starter-parent` POM you can simply declare the plugin and your jars will be repackaged:

```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
    </plugin>
  </plugins>
</build>

```

If you are not using the parent POM you can still use the plugin, however, you must additionally add an `<executions>` section:

```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>

```

```

<artifactId>spring-boot-maven-plugin</artifactId>
<version>2.0.0.M3</version>
<executions>
  <execution>
    <goals>
      <goal>repackage</goal>
    </goals>
  </execution>
</executions>
</plugin>
</plugins>
</build>

```

See the [plugin documentation](#) for full usage details.

85.5 Use a Spring Boot application as a dependency

Like a war file, a Spring Boot application is not intended to be used as a dependency. If your application contains classes that you want to share with other projects, the recommended approach is to move that code into a separate module. The separate module can then be depended upon by your application and other projects.

If you cannot rearrange your code as recommended above, Spring Boot's Maven and Gradle plugins must be configured to produce a separate artifact that is suitable for use as a dependency. The executable archive cannot be used as a dependency as the [executable jar format](#) packages application classes in `BOOT-INF/classes`. This means that they cannot be found when the executable jar is used as a dependency.

To produce the two artifacts, one that can be used as a dependency and one that is executable, a classifier must be specified. This classifier is applied to the name of the executable archive, leaving the default archive for use as dependency.

To configure a classifier of `exec` in Maven, the following configuration can be used:

```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
      <configuration>
        <classifier>exec</classifier>
      </configuration>
    </plugin>
  </plugins>
</build>

```

85.6 Extract specific libraries when an executable jar runs

Most nested libraries in an executable jar do not need to be unpacked in order to run, however, certain libraries can have problems. For example, JRuby includes its own nested jar support which assumes that the `jruby-complete.jar` is always directly available as a file in its own right.

To deal with any problematic libraries, you can flag that specific nested jars should be automatically unpacked to the 'temp folder' when the executable jar first runs.

For example, to indicate that JRuby should be flagged for unpack using the Maven Plugin you would add the following configuration:

```

<build>

```

```

<plugins>
  <plugin>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-maven-plugin</artifactId>
    <configuration>
      <requiresUnpack>
        <dependency>
          <groupId>org.jruby</groupId>
          <artifactId>jruby-complete</artifactId>
        </dependency>
      </requiresUnpack>
    </configuration>
  </plugin>
</plugins>
</build>

```

85.7 Create a non-executable JAR with exclusions

Often if you have an executable and a non-executable jar as build products, the executable version will have additional configuration files that are not needed in a library jar. E.g. the `application.yml` configuration file might be excluded from the non-executable JAR.

The `maven-jar-plugin` is used to expose a `forceCreation` attribute that allows you to create the jar *again* once the `repackage` goal has run. Arguably, this was a bit fragile anyway since it was relying on the order of plugin executions. In Maven, the executable jar must be the main artifact and you can add a classified jar for the library:

```

<build>
  <plugins>
    <plugin>
      <groupId>org.springframework.boot</groupId>
      <artifactId>spring-boot-maven-plugin</artifactId>
    </plugin>
    <plugin>
      <artifactId>maven-jar-plugin</artifactId>
      <executions>
        <execution>
          <id>lib</id>
          <phase>package</phase>
          <goals>
            <goal>jar</goal>
          </goals>
          <configuration>
            <classifier>lib</classifier>
            <excludes>
              <exclude>application.yml</exclude>
            </excludes>
          </configuration>
        </execution>
      </executions>
    </plugin>
  </plugins>
</build>

```

85.8 Remote debug a Spring Boot application started with Maven

To attach a remote debugger to a Spring Boot application started with Maven you can use the `jvmArguments` property of the [maven plugin](#).

Check [this example](#) for more details.

85.9 Build an executable archive from Ant without using spring-boot-antlib

To build with Ant you need to grab dependencies, compile and then create a jar or war archive. To make it executable you can either use the `spring-boot-antlib` module, or you can follow these instructions:

1. If you are building a jar, package the application's classes and resources in a nested `BOOT-INF/classes` directory. If you are building a war, package the application's classes in a nested `WEB-INF/classes` directory as usual.
2. Add the runtime dependencies in a nested `BOOT-INF/lib` directory for a jar or `WEB-INF/lib` for a war. Remember **not** to compress the entries in the archive.
3. Add the `provided` (embedded container) dependencies in a nested `BOOT-INF/lib` directory for jar or `WEB-INF/lib-provided` for a war. Remember **not** to compress the entries in the archive.
4. Add the `spring-boot-loader` classes at the root of the archive (so the `Main-Class` is available).
5. Use the appropriate launcher, e.g. `JarLauncher` for a jar file, as a `Main-Class` attribute in the manifest and specify the other properties it needs as manifest entries, principally a `Start-Class`.

Example:

```
<target name="build" depends="compile">
  <jar destfile="target/${ant.project.name}-${spring-boot.version}.jar" compress="false">
    <mappedresources>
      <fileset dir="target/classes" />
      <globmapper from="*" to="BOOT-INF/classes/*"/>
    </mappedresources>
    <mappedresources>
      <fileset dir="src/main/resources" erroronmissingdir="false"/>
      <globmapper from="*" to="BOOT-INF/classes/*"/>
    </mappedresources>
    <mappedresources>
      <fileset dir="${lib.dir}/runtime" />
      <globmapper from="*" to="BOOT-INF/lib/*"/>
    </mappedresources>
    <zipfileset src="${lib.dir}/loader/spring-boot-loader-jar-${spring-boot.version}.jar" />
    <manifest>
      <attribute name="Main-Class" value="org.springframework.boot.loader.JarLauncher" />
      <attribute name="Start-Class" value="${start-class}" />
    </manifest>
  </jar>
</target>
```

The [Ant Sample](#) has a `build.xml` with a manual task that should work if you run it with

```
$ ant -lib <folder containing ivy-2.2.jar> clean manual
```

after which you can run the application with

```
$ java -jar target/*.jar
```

86. Traditional deployment

86.1 Create a deployable war file

The first step in producing a deployable war file is to provide a `SpringBootServletInitializer` subclass and override its `configure` method. This makes use of Spring Framework's Servlet 3.0 support and allows you to configure your application when it's launched by the servlet container. Typically, you update your application's main class to extend `SpringBootServletInitializer`:

```
@SpringBootApplication
public class Application extends SpringBootServletInitializer {

    @Override
    protected SpringApplicationBuilder configure(SpringApplicationBuilder application) {
        return application.sources(Application.class);
    }

    public static void main(String[] args) throws Exception {
        SpringApplication.run(Application.class, args);
    }
}
```

The next step is to update your build configuration so that your project produces a war file rather than a jar file. If you're using Maven and using `spring-boot-starter-parent` (which configures Maven's war plugin for you) all you need to do is to modify `pom.xml` to change the packaging to war:

```
<packaging>war</packaging>
```

If you're using Gradle, you need to modify `build.gradle` to apply the war plugin to the project:

```
apply plugin: 'war'
```

The final step in the process is to ensure that the embedded servlet container doesn't interfere with the servlet container to which the war file will be deployed. To do so, you need to mark the embedded servlet container dependency as provided.

If you're using Maven:

```
<dependencies>
  <!-- ... -->
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-tomcat</artifactId>
    <scope>provided</scope>
  </dependency>
  <!-- ... -->
</dependencies>
```

And if you're using Gradle:

```
dependencies {
    // ...
    providedRuntime 'org.springframework.boot:spring-boot-starter-tomcat'
    // ...
}
```

Note

`providedRuntime` is preferred to Gradle's `compileOnly` configuration as, among other limitations, `compileOnly` dependencies are not on the test classpath so any web-based integration tests will fail.

If you're using the [Spring Boot build tools](#), marking the embedded servlet container dependency as `provided` will produce an executable war file with the provided dependencies packaged in a `lib-provided` directory. This means that, in addition to being deployable to a servlet container, you can also run your application using `java -jar` on the command line.

Tip

Take a look at Spring Boot's sample applications for a [Maven-based example](#) of the above-described configuration.

86.2 Create a deployable war file for older servlet containers

Older Servlet containers don't have support for the `ServletContextInitializer` bootstrap process used in Servlet 3.0. You can still use Spring and Spring Boot in these containers but you are going to need to add a `web.xml` to your application and configure it to load an `ApplicationContext` via a `DispatcherServlet`.

86.3 Convert an existing application to Spring Boot

For a non-web application it should be easy (throw away the code that creates your `ApplicationContext` and replace it with calls to `SpringApplication` or `SpringApplicationBuilder`). Spring MVC web applications are generally amenable to first creating a deployable war application, and then migrating it later to an executable war and/or jar. Useful reading is in the [Getting Started Guide on Converting a jar to a war](#).

Create a deployable war by extending `SpringBootServletInitializer` (e.g. in a class called `Application`), and add the Spring Boot `@SpringBootApplication` annotation. Example:

```
@SpringBootApplication
public class Application extends SpringBootServletInitializer {

    @Override
    protected SpringApplicationBuilder configure(SpringApplicationBuilder application) {
        // Customize the application or call application.sources(...) to add sources
        // Since our example is itself a @Configuration class (via @SpringBootApplication)
        // we actually don't need to override this method.
        return application;
    }
}
```

Remember that whatever you put in the `sources` is just a Spring `ApplicationContext` and normally anything that already works should work here. There might be some beans you can remove later and let Spring Boot provide its own defaults for them, but it should be possible to get something working first.

Static resources can be moved to `/public` (or `/static` or `/resources` or `/META-INF/resources`) in the classpath root. Same for `messages.properties` (Spring Boot detects this automatically in the root of the classpath).

Vanilla usage of `SpringDispatcherServlet` and `Spring Security` should require no further changes. If you have other features in your application, using other servlets or filters for instance, then you may need to add some configuration to your `ApplicationContext`, replacing those elements from the `web.xml` as follows:

- A `@Bean` of type `Servlet` or `ServletRegistrationBean` installs that bean in the container as if it was a `<servlet/>` and `<servlet-mapping/>` in `web.xml`.
- A `@Bean` of type `Filter` or `FilterRegistrationBean` behaves similarly (like a `<filter/>` and `<filter-mapping/>`).
- An `ApplicationContext` in an XML file can be added through an `@ImportResource` in your `Application`. Or simple cases where annotation configuration is heavily used already can be recreated in a few lines as `@Bean` definitions.

Once the war is working we make it executable by adding a `main` method to our `Application`, e.g.

```
public static void main(String[] args) {
    SpringApplication.run(Application.class, args);
}
```

Note

If you intend to start your application as a war or as an executable application, you need to share the customizations of the builder in a method that is both available to the `SpringBootServletInitializer` callback and the `main` method, something like:

```
@SpringBootApplication
public class Application extends SpringBootServletInitializer {

    @Override
    protected SpringApplicationBuilder configure(SpringApplicationBuilder builder) {
        return configureApplication(builder);
    }

    public static void main(String[] args) {
        configureApplication(new SpringApplicationBuilder()).run(args);
    }

    private static SpringApplicationBuilder configureApplication(SpringApplicationBuilder
builder) {
        return builder.sources(Application.class).bannerMode(Banner.Mode.OFF);
    }
}
```

Applications can fall into more than one category:

- Servlet 3.0+ applications with no `web.xml`.
- Applications with a `web.xml`.
- Applications with a context hierarchy.
- Applications without a context hierarchy.

All of these should be amenable to translation, but each might require slightly different tricks.

Servlet 3.0+ applications might translate pretty easily if they already use the `Spring Servlet 3.0+` initializer support classes. Normally all the code from an existing `WebApplicationInitializer`

can be moved into a `SpringBootServletInitializer`. If your existing application has more than one `ApplicationContext` (e.g. if it uses `AbstractDispatcherServletInitializer`) then you might be able to squash all your context sources into a single `SpringApplication`. The main complication you might encounter is if that doesn't work and you need to maintain the context hierarchy. See the [entry on building a hierarchy](#) for examples. An existing parent context that contains web-specific features will usually need to be broken up so that all the `ServletContextAware` components are in the child context.

Applications that are not already Spring applications might be convertible to a Spring Boot application, and the guidance above might help, but your mileage may vary.

86.4 Deploying a WAR to WebLogic

To deploy a Spring Boot application to WebLogic you must ensure that your servlet initializer **directly** implements `WebApplicationInitializer` (even if you extend from a base class that already implements it).

A typical initializer for WebLogic would be something like this:

```
import org.springframework.boot.autoconfigure.SpringBootApplication;
import org.springframework.boot.context.web.SpringBootServletInitializer;
import org.springframework.web.WebApplicationInitializer;

@SpringBootApplication
public class MyApplication extends SpringBootServletInitializer implements WebApplicationInitializer {
}

```

If you use logback, you will also need to tell WebLogic to prefer the packaged version rather than the version that pre-installed with the server. You can do this by adding a `WEB-INF/weblogic.xml` file with the following contents:

```
<?xml version="1.0" encoding="UTF-8"?>
<wls:weblogic-web-app
  xmlns:wls="http://xmlns.oracle.com/weblogic/weblogic-web-app"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
    http://java.sun.com/xml/ns/javaee/ejb-jar_3_0.xsd
    http://xmlns.oracle.com/weblogic/weblogic-web-app
    http://xmlns.oracle.com/weblogic/weblogic-web-app/1.4/weblogic-web-app.xsd">
  <wls:container-descriptor>
    <wls:prefer-application-packages>
      <wls:package-name>org.slf4j</wls:package-name>
    </wls:prefer-application-packages>
  </wls:container-descriptor>
</wls:weblogic-web-app>

```

86.5 Deploying a WAR in an Old (Servlet 2.5) Container

Spring Boot uses Servlet 3.0 APIs to initialize the `ServletContext` (register Servlets etc.) so you can't use the same application out of the box in a Servlet 2.5 container. It **is** however possible to run a Spring Boot application on an older container with some special tools. If you include `org.springframework.boot:spring-boot-legacy` as a dependency ([maintained separately](#) to the core of Spring Boot and currently available at 1.0.2.RELEASE), all you should need to do is create a `web.xml` and declare a context listener to create the application context and your filters and servlets. The context listener is a special purpose one for Spring Boot, but the rest of it is normal for a Spring application in Servlet 2.5. Example:

```

<?xml version="1.0" encoding="UTF-8"?>
<web-app version="2.5" xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee http://java.sun.com/xml/ns/javaee/web-
app_2_5.xsd">

  <context-param>
    <param-name>contextConfigLocation</param-name>
    <param-value>demo.Application</param-value>
  </context-param>

  <listener>
    <listener-class>org.springframework.boot.legacy.context.web.SpringBootContextLoaderListener</
listener-class>
  </listener>

  <filter>
    <filter-name>metricsFilter</filter-name>
    <filter-class>org.springframework.web.filter.DelegatingFilterProxy</filter-class>
  </filter>

  <filter-mapping>
    <filter-name>metricsFilter</filter-name>
    <url-pattern>/*</url-pattern>
  </filter-mapping>

  <servlet>
    <servlet-name>appServlet</servlet-name>
    <servlet-class>org.springframework.web.servlet.DispatcherServlet</servlet-class>
    <init-param>
      <param-name>contextAttribute</param-name>
      <param-value>org.springframework.web.context.WebApplicationContext.ROOT</param-value>
    </init-param>
    <load-on-startup>1</load-on-startup>
  </servlet>

  <servlet-mapping>
    <servlet-name>appServlet</servlet-name>
    <url-pattern>/</url-pattern>
  </servlet-mapping>

</web-app>

```

In this example we are using a single application context (the one created by the context listener) and attaching it to the `DispatcherServlet` using an init parameter. This is normal in a Spring Boot application (you normally only have one application context).

86.6 Use Lettuce instead of Jedis

The Spring Boot starter (`spring-boot-starter-data-redis`) uses [Jedis](#) by default. You need to exclude that dependency and include the [Lettuce](#) one instead. You also need `commons-pool2`. Spring Boot manages these dependencies to help make this process as easy as possible.

Example in Maven:

```

<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-data-redis</artifactId>
  <exclusions>
    <exclusion>
      <groupId>redis.clients</groupId>
      <artifactId>jedis</artifactId>
    </exclusion>
  </exclusions>
</dependency>
</dependency>

```

```
<groupId>io.lettuce</groupId>
<artifactId>lettuce-core</artifactId>
</dependency>
<dependency>
  <groupId>org.apache.commons</groupId>
  <artifactId>commons-pool2</artifactId>
</dependency>
```

Example in Gradle:

```
configurations {
  compile.exclude module: "jedis"
}

dependencies {
  compile("io.lettuce:lettuce-core")
  compile("org.apache.commons:commons-pool2")
  // ...
}
```

Part X. Appendices

Appendix A. Common application properties

Various properties can be specified inside your `application.properties/application.yml` file or as command line switches. This section provides a list of common Spring Boot properties and references to the underlying classes that consume them.

Note

Property contributions can come from additional jar files on your classpath so you should not consider this an exhaustive list. It is also perfectly legit to define your own properties.

Warning

This sample file is meant as a guide only. Do **not** copy/paste the entire content into your application; rather pick only the properties that you need.

```
# =====
# COMMON SPRING BOOT PROPERTIES
#
# This sample file is provided as a guideline. Do NOT copy it in its
# entirety to your own application.    ^^^
# =====

# -----
# CORE PROPERTIES
# -----

# BANNER
banner.charset=UTF-8 # Banner file encoding.
banner.location=classpath:banner.txt # Banner file location.
banner.image.location=classpath:banner.gif # Banner image file location (jpg/png can also be used).
banner.image.width= # Width of the banner image in chars (default 76)
banner.image.height= # Height of the banner image in chars (default based on image height)
banner.image.margin= # Left hand image margin in chars (default 2)
banner.image.invert= # If images should be inverted for dark terminal themes (default false)

# LOGGING
logging.config= # Location of the logging configuration file. For instance `classpath:logback.xml` for
Logback
logging.exception-conversion-word=%wEx # Conversion word used when logging exceptions.
logging.file= # Log file name. For instance `myapp.log`
logging.level.*= # Log levels severity mapping. For instance `logging.level.org.springframework=DEBUG`
logging.path= # Location of the log file. For instance `/var/log`
logging.pattern.console= # Appender pattern for output to the console. Only supported with the default
logback setup.
logging.pattern.file= # Appender pattern for output to the file. Only supported with the default logback
setup.
logging.pattern.level= # Appender pattern for log level (default %5p). Only supported with the default
logback setup.
logging.register-shutdown-hook=false # Register a shutdown hook for the logging system when it is
initialized.

# AOP
spring.aop.auto=true # Add @EnableAspectJAutoProxy.
spring.aop.proxy-target-class=true # Whether subclass-based (CGLIB) proxies are to be created (true) as
opposed to standard Java interface-based proxies (false).

# IDENTITY (ContextIdApplicationContextInitializer)
spring.application.index= # Application index.
```

```

spring.application.name= # Application name.

# ADMIN (SpringApplicationAdminJmxAutoConfiguration)
spring.application.admin.enabled=false # Enable admin features for the application.
spring.application.admin.jmx-name=org.springframework.boot:type=Admin,name=SpringApplication # JMX name
of the application admin MBean.

# AUTO-CONFIGURATION
spring.autoconfigure.exclude= # Auto-configuration classes to exclude.

# SPRING CORE
spring.beaninfo.ignore=true # Skip search of BeanInfo classes.

# SPRING CACHE (CacheProperties)
spring.cache.cache-names= # Comma-separated list of cache names to create if supported by the underlying
cache manager.
spring.cache.caffeine.spec= # The spec to use to create caches. Check CaffeineSpec for more details on
the spec format.
spring.cache.couchbase.expiration=0 # Entry expiration in milliseconds. By default the entries never
expire.
spring.cache.ehcache.config= # The location of the configuration file to use to initialize EhCache.
spring.cache.infinispan.config= # The location of the configuration file to use to initialize
Infinispan.
spring.cache.jcache.config= # The location of the configuration file to use to initialize the cache
manager.
spring.cache.jcache.provider= # Fully qualified name of the CachingProvider implementation to use to
retrieve the JSR-107 compliant cache manager. Only needed if more than one JSR-107 implementation is
available on the classpath.
spring.cache.type= # Cache type, auto-detected according to the environment by default.

# SPRING CONFIG - using environment property only (ConfigFileApplicationListener)
spring.config.location= # Config file locations.
spring.config.name=application # Config file name.

# HAZELCAST (HazelcastProperties)
spring.hazelcast.config= # The location of the configuration file to use to initialize Hazelcast.

# PROJECT INFORMATION (ProjectInfoProperties)
spring.info.build.location=classpath:META-INF/build-info.properties # Location of the generated build-
info.properties file.
spring.info.git.location=classpath:git.properties # Location of the generated git.properties file.

# JMX
spring.jmx.default-domain= # JMX domain name.
spring.jmx.enabled=true # Expose management beans to the JMX domain.
spring.jmx.server=mbeanServer # MBeanServer bean name.

# Email (MailProperties)
spring.mail.default-encoding=UTF-8 # Default MimeMessage encoding.
spring.mail.host= # SMTP server host. For instance `smtp.example.com`
spring.mail.jndi-name= # Session JNDI name. When set, takes precedence to others mail settings.
spring.mail.password= # Login password of the SMTP server.
spring.mail.port= # SMTP server port.
spring.mail.properties.*= # Additional JavaMail session properties.
spring.mail.protocol=smtp # Protocol used by the SMTP server.
spring.mail.test-connection=false # Test that the mail server is available on startup.
spring.mail.username= # Login user of the SMTP server.

# APPLICATION SETTINGS (SpringApplication)
spring.main.banner-mode=console # Mode used to display the banner when the application runs.
spring.main.sources= # Sources (class name, package name or XML resource location) to include in the
ApplicationContext.
spring.main.web-application-type= # Flag to explicitly request a specific type of web application. Auto-
detected based on the classpath if not set.

# FILE ENCODING (FileEncodingApplicationListener)
spring.mandatory-file-encoding= # Expected character encoding the application must use.

# INTERNATIONALIZATION (MessageSourceAutoConfiguration)
spring.messages.always-use-message-format=false # Set whether to always apply the MessageFormat rules,
parsing even messages without arguments.

```

```

spring.messages.basename=messages # Comma-separated list of basenames, each following the ResourceBundle
convention.
spring.messages.cache-seconds=-1 # Loaded resource bundle files cache expiration, in seconds. When set
to -1, bundles are cached forever.
spring.messages.encoding=UTF-8 # Message bundles encoding.
spring.messages.fallback-to-system-locale=true # Set whether to fall back to the system Locale if no
files for a specific Locale have been found.

# OUTPUT
spring.output.ansi.enabled=detect # Configure the ANSI output.

# PID FILE (ApplicationPidFileWriter)
spring.pid.fail-on-write-error= # Fail if ApplicationPidFileWriter is used but it cannot write the PID
file.
spring.pid.file= # Location of the PID file to write (if ApplicationPidFileWriter is used).

# PROFILES
spring.profiles.active= # Comma-separated list (or list if using YAML) of active profiles.
spring.profiles.include= # Unconditionally activate the specified comma separated profiles (or list of
profiles if using YAML).

# QUARTZ SCHEDULER (QuartzProperties)
spring.quartz.job-store-type=memory # Quartz job store type.
spring.quartz.properties.*= # Additional Quartz Scheduler properties.
spring.quartz.jdbc.initialize-schema=false # Create the required Quartz Scheduler tables on startup.
spring.quartz.jdbc.schema=classpath:org/quartz/impl/jdbcjobstore/tables_@@platform@@.sql # Path to the
SQL file to use to initialize the database schema.

# Reactor
spring.reactor.stacktrace-mode.enabled=false # Set whether Reactor should collect stacktrace information
at runtime.

# SENDGRID (SendGridAutoConfiguration)
spring.sendgrid.api-key= # SendGrid API key.
spring.sendgrid.proxy.host= # SendGrid proxy host.
spring.sendgrid.proxy.port= # SendGrid proxy port.

# -----
# WEB PROPERTIES
# -----

# EMBEDDED SERVER CONFIGURATION (ServerProperties)
server.address= # Network address to which the server should bind to.
server.compression.enabled=false # If response compression is enabled.
server.compression.excluded-user-agents= # List of user-agents to exclude from compression.
server.compression.mime-types= # Comma-separated list of MIME types that should be compressed. For
instance `text/html,text/css,application/json`
server.compression.min-response-size= # Minimum response size that is required for compression to be
performed. For instance 2048
server.connection-timeout= # Time in milliseconds that connectors will wait for another HTTP request
before closing the connection. When not set, the connector's container-specific default will be used.
Use a value of -1 to indicate no (i.e. infinite) timeout.
server.display-name=application # Display name of the application.
server.max-http-header-size=0 # Maximum size in bytes of the HTTP message header.
server.error.include-exception=false # Include the "exception" attribute.
server.error.include-stacktrace=never # When to include a "stacktrace" attribute.
server.error.path=/error # Path of the error controller.
server.error.whitelabel.enabled=true # Enable the default error page displayed in browsers in case of a
server error.
server.jetty.acceptors= # Number of acceptor threads to use.
server.jetty.accesslog.append=false # Append to log.
server.jetty.accesslog.date-format=dd/MMM/yyyy:HH:mm:ss Z # Timestamp format of the request log.
server.jetty.accesslog.enabled=false # Enable access log.
server.jetty.accesslog.extended-format=false # Enable extended NCSA format.
server.jetty.accesslog.file-date-format= # Date format to place in log file name.
server.jetty.accesslog.filename= # Log filename. If not specified, logs will be redirected to
"System.err".
server.jetty.accesslog.locale= # Locale of the request log.
server.jetty.accesslog.log-cookies=false # Enable logging of the request cookies.
server.jetty.accesslog.log-latency=false # Enable logging of request processing time.

```

```

server.jetty.accesslog.log-server=false # Enable logging of the request hostname.
server.jetty.accesslog.retention-period=31 # Number of days before rotated log files are deleted.
server.jetty.accesslog.time-zone=GMT # Timezone of the request log.
server.jetty.max-http-post-size=0 # Maximum size in bytes of the HTTP post or put content.
server.jetty.selectors= # Number of selector threads to use.
server.port=8080 # Server HTTP port.
server.server-header= # Value to use for the Server response header (no header is sent if empty)
server.use-forward-headers= # If X-Forwarded-* headers should be applied to the HttpRequest.
server.servlet.context-parameters.*= # Servlet context init parameters
server.servlet.context-path= # Context path of the application.
server.servlet.jsp.class-name=org.apache.jasper.servlet.JspServlet # The class name of the JSP servlet.
server.servlet.jsp.init-parameters.*= # Init parameters used to configure the JSP servlet
server.servlet.jsp.registered=true # Whether or not the JSP servlet is registered
server.servlet.path=/ # Path of the main dispatcher servlet.
server.session.cookie.comment= # Comment for the session cookie.
server.session.cookie.domain= # Domain for the session cookie.
server.session.cookie.http-only= # "HttpOnly" flag for the session cookie.
server.session.cookie.max-age= # Maximum age of the session cookie in seconds.
server.session.cookie.name= # Session cookie name.
server.session.cookie.path= # Path of the session cookie.
server.session.cookie.secure= # "Secure" flag for the session cookie.
server.session.persistent=false # Persist session data between restarts.
server.session.servlet.filter-order=-2147483598 # Session repository filter order.
server.session.servlet.filter-dispatcher-types=ASYNC, ERROR, REQUEST # Session repository filter
dispatcher types.
server.session.store-dir= # Directory used to store session data.
server.session.timeout= # Session timeout in seconds.
server.session.tracking-modes= # Session tracking modes (one or more of the following: "cookie", "url",
"ssl").
server.ssl.ciphers= # Supported SSL ciphers.
server.ssl.client-auth= # Whether client authentication is wanted ("want") or needed ("need"). Requires
a trust store.
server.ssl.enabled= # Enable SSL support.
server.ssl.enabled-protocols= # Enabled SSL protocols.
server.ssl.key-alias= # Alias that identifies the key in the key store.
server.ssl.key-password= # Password used to access the key in the key store.
server.ssl.key-store= # Path to the key store that holds the SSL certificate (typically a jks file).
server.ssl.key-store-password= # Password used to access the key store.
server.ssl.key-store-provider= # Provider for the key store.
server.ssl.key-store-type= # Type of the key store.
server.ssl.protocol=TLS # SSL protocol to use.
server.ssl.trust-store= # Trust store that holds SSL certificates.
server.ssl.trust-store-password= # Password used to access the trust store.
server.ssl.trust-store-provider= # Provider for the trust store.
server.ssl.trust-store-type= # Type of the trust store.
server.tomcat.accept-count= # Maximum queue length for incoming connection requests when all possible
request processing threads are in use.
server.tomcat.accesslog.buffered=true # Buffer output such that it is only flushed periodically.
server.tomcat.accesslog.directory=logs # Directory in which log files are created. Can be relative to
the tomcat base dir or absolute.
server.tomcat.accesslog.enabled=false # Enable access log.
server.tomcat.accesslog.file-date-format=.yyyy-MM-dd # Date format to place in log file name.
server.tomcat.accesslog.pattern=common # Format pattern for access logs.
server.tomcat.accesslog.prefix=access_log # Log file name prefix.
server.tomcat.accesslog.rename-on-rotate=false # Defer inclusion of the date stamp in the file name
until rotate time.
server.tomcat.accesslog.request-attributes-enabled=false # Set request attributes for IP address,
Hostname, protocol and port used for the request.
server.tomcat.accesslog.rotate=true # Enable access log rotation.
server.tomcat.accesslog.suffix=.log # Log file name suffix.
server.tomcat.additional-tld-skip-patterns= # Comma-separated list of additional patterns that match
jars to ignore for TLD scanning.
server.tomcat.background-processor-delay=30 # Delay in seconds between the invocation of
backgroundProcess methods.
server.tomcat.basedir= # Tomcat base directory. If not specified a temporary directory will be used.
server.tomcat.internal-proxies=10\\.\\d{1,3}\\.\\d{1,3}\\.\\d{1,3}|\\
192\\.168\\.\\d{1,3}\\.\\d{1,3}|\\
169\\.254\\.\\d{1,3}\\.\\d{1,3}|\\
127\\.\\d{1,3}\\.\\d{1,3}\\.\\d{1,3}|\\
172\\.1[6-9]{1}\\.\\d{1,3}\\.\\d{1,3}|\\
172\\.2[0-9]{1}\\.\\d{1,3}\\.\\d{1,3}|\\

```

```

172\\.3[0-1]{1}\\\\.\\d{1,3}\\\\.\\d{1,3} # regular expression matching trusted IP addresses.
server.tomcat.max-connections= # Maximum number of connections that the server will accept and process
at any given time.
server.tomcat.max-http-header-size=0 # Maximum size in bytes of the HTTP message header.
server.tomcat.max-http-post-size=0 # Maximum size in bytes of the HTTP post content.
server.tomcat.max-threads=0 # Maximum amount of worker threads.
server.tomcat.min-spare-threads=0 # Minimum amount of worker threads.
server.tomcat.port-header=X-Forwarded-Port # Name of the HTTP header used to override the original port
value.
server.tomcat.protocol-header= # Header that holds the incoming protocol, usually named "X-Forwarded-
Proto".
server.tomcat.protocol-header-https-value=https # Value of the protocol header that indicates that the
incoming request uses SSL.
server.tomcat.redirect-context-root= # Whether requests to the context root should be redirected by
appending a / to the path.
server.tomcat.remote-ip-header= # Name of the http header from which the remote ip is extracted. For
instance `X-FORWARDED-FOR`
server.tomcat.uri-encoding=UTF-8 # Character encoding to use to decode the URI.
server.undertow.accesslog.dir= # Undertow access log directory.
server.undertow.accesslog.enabled=false # Enable access log.
server.undertow.accesslog.pattern=common # Format pattern for access logs.
server.undertow.accesslog.prefix=access_log. # Log file name prefix.
server.undertow.accesslog.rotate=true # Enable access log rotation.
server.undertow.accesslog.suffix=log # Log file name suffix.
server.undertow.buffer-size= # Size of each buffer in bytes.
server.undertow.direct-buffers= # Allocate buffers outside the Java heap.
server.undertow.io-threads= # Number of I/O threads to create for the worker.
server.undertow.eager-filter-init=true # Whether servlet filters should be initialized on startup.
server.undertow.max-http-post-size=0 # Maximum size in bytes of the HTTP post content.
server.undertow.worker-threads= # Number of worker threads.

# FREEMARKER (FreeMarkerAutoConfiguration)
spring.freemarker.allow-request-override=false # Set whether HttpServletRequest attributes are allowed
to override (hide) controller generated model attributes of the same name.
spring.freemarker.allow-session-override=false # Set whether HttpSession attributes are allowed to
override (hide) controller generated model attributes of the same name.
spring.freemarker.cache=false # Enable template caching.
spring.freemarker.charset=UTF-8 # Template encoding.
spring.freemarker.check-template-location=true # Check that the templates location exists.
spring.freemarker.content-type=text/html # Content-Type value.
spring.freemarker.enabled=true # Enable MVC view resolution for this technology.
spring.freemarker.expose-request-attributes=false # Set whether all request attributes should be added
to the model prior to merging with the template.
spring.freemarker.expose-session-attributes=false # Set whether all HttpSession attributes should be
added to the model prior to merging with the template.
spring.freemarker.expose-spring-macro-helpers=true # Set whether to expose a RequestContext for use by
Spring's macro library, under the name "springMacroRequestContext".
spring.freemarker.prefer-file-system-access=true # Prefer file system access for template loading. File
system access enables hot detection of template changes.
spring.freemarker.prefix= # Prefix that gets prepended to view names when building a URL.
spring.freemarker.request-context-attribute= # Name of the RequestContext attribute for all views.
spring.freemarker.settings.*= # Well-known FreeMarker keys which will be passed to FreeMarker's
Configuration.
spring.freemarker.suffix= # Suffix that gets appended to view names when building a URL.
spring.freemarker.template-loader-path=classpath:/templates/ # Comma-separated list of template paths.
spring.freemarker.view-names= # White list of view names that can be resolved.

# GROOVY TEMPLATES (GroovyTemplateAutoConfiguration)
spring.groovy.template.allow-request-override=false # Set whether HttpServletRequest attributes are
allowed to override (hide) controller generated model attributes of the same name.
spring.groovy.template.allow-session-override=false # Set whether HttpSession attributes are allowed to
override (hide) controller generated model attributes of the same name.
spring.groovy.template.cache= # Enable template caching.
spring.groovy.template.charset=UTF-8 # Template encoding.
spring.groovy.template.check-template-location=true # Check that the templates location exists.
spring.groovy.template.configuration.*= # See GroovyMarkupConfigurer
spring.groovy.template.content-type=test/html # Content-Type value.
spring.groovy.template.enabled=true # Enable MVC view resolution for this technology.
spring.groovy.template.expose-request-attributes=false # Set whether all request attributes should be
added to the model prior to merging with the template.

```

```

spring.groovy.template.expose-session-attributes=false # Set whether all HttpSession attributes should
be added to the model prior to merging with the template.
spring.groovy.template.expose-spring-macro-helpers=true # Set whether to expose a RequestContext for use
by Spring's macro library, under the name "springMacroRequestContext".
spring.groovy.template.prefix= # Prefix that gets prepended to view names when building a URL.
spring.groovy.template.request-context-attribute= # Name of the RequestContext attribute for all views.
spring.groovy.template.resource-loader-path=classpath:/templates/ # Template path.
spring.groovy.template.suffix=.tpl # Suffix that gets appended to view names when building a URL.
spring.groovy.template.view-names= # White list of view names that can be resolved.

# SPRING HATEOAS (HateoasProperties)
spring.hateoas.use-hal-as-default-json-media-type=true # Specify if application/hal+json responses
should be sent to requests that accept application/json.

# HTTP message conversion
spring.http.converters.preferred-json-mapper=jackson # Preferred JSON mapper to use for HTTP message
conversion. Set to "gson" to force the use of Gson when both it and Jackson are on the classpath.

# HTTP encoding (HttpEncodingProperties)
spring.http.encoding.charset=UTF-8 # Charset of HTTP requests and responses. Added to the "Content-Type"
header if not set explicitly.
spring.http.encoding.enabled=true # Enable http encoding support.
spring.http.encoding.force= # Force the encoding to the configured charset on HTTP requests and
responses.
spring.http.encoding.force-request= # Force the encoding to the configured charset on HTTP requests.
Defaults to true when "force" has not been specified.
spring.http.encoding.force-response= # Force the encoding to the configured charset on HTTP responses.
spring.http.encoding.mapping= # Locale to Encoding mapping.

# MULTIPART (MultipartProperties)
spring.servlet.multipart.enabled=true # Enable support of multipart uploads.
spring.servlet.multipart.file-size-threshold=0 # Threshold after which files will be written to disk.
Values can use the suffixes "MB" or "KB" to indicate megabytes or kilobytes respectively.
spring.servlet.multipart.location= # Intermediate location of uploaded files.
spring.servlet.multipart.max-file-size=1MB # Max file size. Values can use the suffixes "MB" or "KB" to
indicate megabytes or kilobytes respectively.
spring.servlet.multipart.max-request-size=10MB # Max request size. Values can use the suffixes "MB" or
"KB" to indicate megabytes or kilobytes respectively.
spring.servlet.multipart.resolve-lazily=false # Whether to resolve the multipart request lazily at the
time of file or parameter access.

# JACKSON (JacksonProperties)
spring.jackson.date-format= # Date format string or a fully-qualified date format class name. For
instance `yyyy-MM-dd HH:mm:ss`.
spring.jackson.default-property-inclusion= # Controls the inclusion of properties during serialization.
spring.jackson.deserialization.*= # Jackson on/off features that affect the way Java objects are
deserialized.
spring.jackson.generator.*= # Jackson on/off features for generators.
spring.jackson.joda-date-time-format= # Joda date time format string. If not configured, "date-format"
will be used as a fallback if it is configured with a format string.
spring.jackson.locale= # Locale used for formatting.
spring.jackson.mapper.*= # Jackson general purpose on/off features.
spring.jackson.parser.*= # Jackson on/off features for parsers.
spring.jackson.property-naming-strategy= # One of the constants on Jackson's PropertyNamingStrategy. Can
also be a fully-qualified class name of a PropertyNamingStrategy subclass.
spring.jackson.serialization.*= # Jackson on/off features that affect the way Java objects are
serialized.
spring.jackson.time-zone= # Time zone used when formatting dates. For instance `America/Los_Angeles`

# JERSEY (JerseyProperties)
spring.jersey.application-path= # Path that serves as the base URI for the application. Overrides the
value of "@ApplicationPath" if specified.
spring.jersey.filter.order=0 # Jersey filter chain order.
spring.jersey.init.*= # Init parameters to pass to Jersey via the servlet or filter.
spring.jersey.servlet.load-on-startup=-1 # Load on startup priority of the Jersey servlet.
spring.jersey.type=servlet # Jersey integration type.

# SPRING LDAP (LdapProperties)
spring.ldap.urls= # LDAP URLs of the server.
spring.ldap.base= # Base suffix from which all operations should originate.
spring.ldap.username= # Login user of the server.

```

```

spring.ldap.password= # Login password of the server.
spring.ldap.base-environment.*= # LDAP specification settings.

# EMBEDDED LDAP (EmbeddedLdapProperties)
spring.ldap.embedded.base-dn= # The base DN
spring.ldap.embedded.credential.username= # Embedded LDAP username.
spring.ldap.embedded.credential.password= # Embedded LDAP password.
spring.ldap.embedded.ldif=classpath:schema.ldif # Schema (LDIF) script resource reference.
spring.ldap.embedded.port= # Embedded LDAP port.
spring.ldap.embedded.validation.enabled=true # Enable LDAP schema validation.
spring.ldap.embedded.validation.schema= # Path to the custom schema.

# SPRING MOBILE DEVICE VIEWS (DeviceDelegatingViewResolverAutoConfiguration)
spring.mobile.devicedelegatingviewresolver.enable-fallback=false # Enable support for fallback
resolution.
spring.mobile.devicedelegatingviewresolver.enabled=false # Enable device view resolver.
spring.mobile.devicedelegatingviewresolver.mobile-prefix=mobile/ # Prefix that gets prepended to view
names for mobile devices.
spring.mobile.devicedelegatingviewresolver.mobile-suffix= # Suffix that gets appended to view names for
mobile devices.
spring.mobile.devicedelegatingviewresolver.normal-prefix= # Prefix that gets prepended to view names for
normal devices.
spring.mobile.devicedelegatingviewresolver.normal-suffix= # Suffix that gets appended to view names for
normal devices.
spring.mobile.devicedelegatingviewresolver.tablet-prefix=tablet/ # Prefix that gets prepended to view
names for tablet devices.
spring.mobile.devicedelegatingviewresolver.tablet-suffix= # Suffix that gets appended to view names for
tablet devices.

# SPRING MOBILE SITE PREFERENCE (SitePreferenceAutoConfiguration)
spring.mobile.sitepreference.enabled=true # Enable SitePreferenceHandler.

# MUSTACHE TEMPLATES (MustacheAutoConfiguration)
spring.mustache.allow-request-override= # Set whether HttpServletRequest attributes are allowed to
override (hide) controller generated model attributes of the same name.
spring.mustache.allow-session-override= # Set whether HttpSession attributes are allowed to override
(hide) controller generated model attributes of the same name.
spring.mustache.cache= # Enable template caching.
spring.mustache.charset= # Template encoding.
spring.mustache.check-template-location= # Check that the templates location exists.
spring.mustache.content-type= # Content-Type value.
spring.mustache.enabled= # Enable MVC view resolution for this technology.
spring.mustache.expose-request-attributes= # Set whether all request attributes should be added to the
model prior to merging with the template.
spring.mustache.expose-session-attributes= # Set whether all HttpSession attributes should be added to
the model prior to merging with the template.
spring.mustache.expose-spring-macro-helpers= # Set whether to expose a RequestContext for use by
Spring's macro library, under the name "springMacroRequestContext".
spring.mustache.prefix=classpath:/templates/ # Prefix to apply to template names.
spring.mustache.request-context-attribute= # Name of the RequestContext attribute for all views.
spring.mustache.suffix=.mustache # Suffix to apply to template names.
spring.mustache.view-names= # White list of view names that can be resolved.

# SPRING MVC (WebMvcProperties)
spring.mvc.async.request-timeout= # Amount of time (in milliseconds) before asynchronous request
handling times out.
spring.mvc.date-format= # Date format to use. For instance `dd/MM/yyyy`.
spring.mvc.dispatch-trace-request=false # Dispatch TRACE requests to the FrameworkServlet doService
method.
spring.mvc.dispatch-options-request=true # Dispatch OPTIONS requests to the FrameworkServlet doService
method.
spring.mvc.favicon.enabled=true # Enable resolution of favicon.ico.
spring.mvc.formcontent.putfilter.enabled=true # Enable Spring's HttpPutFormContentFilter.
spring.mvc.ignore-default-model-on-redirect=true # If the content of the "default" model should be
ignored during redirect scenarios.
spring.mvc.locale= # Locale to use. By default, this locale is overridden by the "Accept-Language"
header.
spring.mvc.locale-resolver=accept-header # Define how the locale should be resolved.
spring.mvc.log-resolved-exception=false # Enable warn logging of exceptions resolved by a
"HandlerExceptionResolver".
spring.mvc.media-types.*= # Maps file extensions to media types for content negotiation.

```

```

spring.mvc.message-codes-resolver-format= # Formatting strategy for message codes. For instance
`PREFIX_ERROR_CODE`.
spring.mvc.servlet.load-on-startup=-1 # Load on startup priority of the Spring Web Services servlet.
spring.mvc.static-path-pattern=/** # Path pattern used for static resources.
spring.mvc.throw-exception-if-no-handler-found=false # If a "NoHandlerFoundException" should be thrown
if no Handler was found to process a request.
spring.mvc.view.prefix= # Spring MVC view prefix.
spring.mvc.view.suffix= # Spring MVC view suffix.

# SPRING RESOURCES HANDLING (ResourceProperties)
spring.resources.add-mappings=true # Enable default resource handling.
spring.resources.cache-period= # Cache period for the resources served by the resource handler, in
seconds.
spring.resources.chain.cache=true # Enable caching in the Resource chain.
spring.resources.chain.enabled= # Enable the Spring Resource Handling chain. Disabled by default unless
at least one strategy has been enabled.
spring.resources.chain.gziped=false # Enable resolution of already gziped resources.
spring.resources.chain.html-application-cache=false # Enable HTML5 application cache manifest rewriting.
spring.resources.chain.strategy.content.enabled=false # Enable the content Version Strategy.
spring.resources.chain.strategy.content.paths=/** # Comma-separated list of patterns to apply to the
Version Strategy.
spring.resources.chain.strategy.fixed.enabled=false # Enable the fixed Version Strategy.
spring.resources.chain.strategy.fixed.paths=/** # Comma-separated list of patterns to apply to the
Version Strategy.
spring.resources.chain.strategy.fixed.version= # Version string to use for the Version Strategy.
spring.resources.static-locations=classpath:/META-INF/resources/,classpath:/resources/,classpath:/
static/,classpath:/public/ # Locations of static resources.

# SPRING SESSION (SessionProperties)
spring.session.hazelcast.flush-mode=on-save # Sessions flush mode.
spring.session.hazelcast.map-name=spring:session:sessions # Name of the map used to store sessions.
spring.session.jdbc.initializer.enabled= # Create the required session tables on startup if necessary.
Enabled automatically if the default table name is set or a custom schema is configured.
spring.session.jdbc.schema=classpath:org/springframework/session/jdbc/schema-@@platform@@.sql # Path to
the SQL file to use to initialize the database schema.
spring.session.jdbc.table-name=SPRING_SESSION # Name of database table used to store sessions.
spring.session.redis.flush-mode=on-save # Sessions flush mode.
spring.session.redis.namespace= # Namespace for keys used to store sessions.
spring.session.store-type= # Session store type.

# SPRING SOCIAL (SocialWebAutoConfiguration)
spring.social.auto-connection-views=false # Enable the connection status view for supported providers.

# SPRING SOCIAL FACEBOOK (FacebookAutoConfiguration)
spring.social.facebook.app-id= # your application's Facebook App ID
spring.social.facebook.app-secret= # your application's Facebook App Secret

# SPRING SOCIAL LINKEDIN (LinkedInAutoConfiguration)
spring.social.linkedin.app-id= # your application's LinkedIn App ID
spring.social.linkedin.app-secret= # your application's LinkedIn App Secret

# SPRING SOCIAL TWITTER (TwitterAutoConfiguration)
spring.social.twitter.app-id= # your application's Twitter App ID
spring.social.twitter.app-secret= # your application's Twitter App Secret

# THYMELEAF (ThymeleafAutoConfiguration)
spring.thymeleaf.cache=true # Enable template caching.
spring.thymeleaf.check-template=true # Check that the template exists before rendering it.
spring.thymeleaf.check-template-location=true # Check that the templates location exists.
spring.thymeleaf.enabled=true # Enable Thymeleaf view resolution for Web frameworks.
spring.thymeleaf.encoding=UTF-8 # Template files encoding.
spring.thymeleaf.excluded-view-names= # Comma-separated list of view names that should be excluded from
resolution.
spring.thymeleaf.mode=HTML5 # Template mode to be applied to templates. See also
StandardTemplateModeHandlers.
spring.thymeleaf.prefix=classpath:/templates/ # Prefix that gets prepended to view names when building a
URL.
spring.thymeleaf.reactive.max-chunk-size= # Maximum size of data buffers used for writing to the
response, in bytes.
spring.thymeleaf.reactive.media-types= # Media types supported by the view technology.
spring.thymeleaf.servlet.content-type=text/html # Content-Type value written to HTTP responses.

```



```

spring.thymeleaf.suffix=.html # Suffix that gets appended to view names when building a URL.
spring.thymeleaf.template-resolver-order= # Order of the template resolver in the chain.
spring.thymeleaf.view-names= # Comma-separated list of view names that can be resolved.

# SPRING WEB FLUX (WebFluxProperties)
spring.webflux.static-path-pattern=/** # Path pattern used for static resources.

# SPRING WEB SERVICES (WebServicesProperties)
spring.webservices.path=/services # Path that serves as the base URI for the services.
spring.webservices.servlet.init= # Servlet init parameters to pass to Spring Web Services.
spring.webservices.servlet.load-on-startup=-1 # Load on startup priority of the Spring Web Services
servlet.

# -----
# SECURITY PROPERTIES
# -----
# SECURITY (SecurityProperties)
security.basic.authorize-mode=role # Security authorize mode to apply.
security.basic.enabled=true # Enable basic authentication.
security.basic.path=/** # Comma-separated list of paths to secure.
security.basic.realm=Spring # HTTP basic realm name.
security.enable-csrf=false # Enable Cross Site Request Forgery support.
security.filter-order=0 # Security filter chain order.
security.filter-dispatcher-types=ASYNC, FORWARD, INCLUDE, REQUEST # Security filter chain dispatcher
types.
security.headers.cache=true # Enable cache control HTTP headers.
security.headers.content-security-policy= # Value for content security policy header.
security.headers.content-security-policy-mode=default # Content security policy mode.
security.headers.content-type=true # Enable "X-Content-Type-Options" header.
security.headers.frame=true # Enable "X-Frame-Options" header.
security.headers.hsts=all # HTTP Strict Transport Security (HSTS) mode (none, domain, all).
security.headers.xss=true # Enable cross site scripting (XSS) protection.
security.ignored= # Comma-separated list of paths to exclude from the default secured paths.
security.require-ssl=false # Enable secure channel for all requests.
security.sessions=stateless # Session creation policy (always, never, if_required, stateless).
security.user.name=user # Default user name.
security.user.password= # Password for the default user name. A random password is logged on startup by
default.
security.user.role=USER # Granted roles for the default user name.

# SECURITY OAUTH2 CLIENT (OAuth2ClientProperties)
security.oauth2.client.client-id= # OAuth2 client id.
security.oauth2.client.client-secret= # OAuth2 client secret. A random secret is generated by default

# SECURITY OAUTH2 RESOURCES (ResourceServerProperties)
security.oauth2.resource.filter-order= # The order of the filter chain used to authenticate tokens.
security.oauth2.resource.id= # Identifier of the resource.
security.oauth2.resource.jwt.key-uri= # The URI of the JWT token. Can be set if the value is not
available and the key is public.
security.oauth2.resource.jwt.key-value= # The verification key of the JWT token. Can either be a
symmetric secret or PEM-encoded RSA public key.
security.oauth2.resource.prefer-token-info=true # Use the token info, can be set to false to use the
user info.
security.oauth2.resource.service-id=resource #
security.oauth2.resource.token-info-uri= # URI of the token decoding endpoint.
security.oauth2.resource.token-type= # The token type to send when using the userInfoUri.
security.oauth2.resource.user-info-uri= # URI of the user endpoint.

# SECURITY OAUTH2 SSO (OAuth2SsoProperties)
security.oauth2.sso.filter-order= # Filter order to apply if not providing an explicit
WebSecurityConfigurerAdapter
security.oauth2.sso.login-path=/login # Path to the login page, i.e. the one that triggers the redirect
to the OAuth2 Authorization Server

# -----
# DATA PROPERTIES
# -----

```

```

# FLYWAY (FlywayProperties)
flyway.allow-mixed-migrations= #
flyway.baseline-description= #
flyway.baseline-on-migrate= #
flyway.baseline-version=1 # version to start migration
flyway.check-location=false # Check that migration scripts location exists.
flyway.clean-disabled= #
flyway.clean-on-validation-error= #
flyway.enabled=true # Enable flyway.
flyway.encoding= #
flyway.ignore-failed-future-migration= #
flyway.ignore-future-migrations= #
flyway.ignore-missing-migrations= #
flyway.init-sqls= # SQL statements to execute to initialize a connection immediately after obtaining it.
flyway.installed-by= #
flyway.locations=classpath:db/migration # locations of migrations scripts
flyway.out-of-order= #
flyway.password= # JDBC password if you want Flyway to create its own DataSource
flyway.placeholder-prefix= #
flyway.placeholder-replacement= #
flyway.placeholder-suffix= #
flyway.placeholders.*= #
flyway.repeatable-sql-migration-prefix= #
flyway.schemas= # schemas to update
flyway.skip-default-callbacks= #
flyway.skip-default-resolvers= #
flyway.sql-migration-prefix=V #
flyway.sql-migration-separator= #
flyway.sql-migration-suffix=.sql #
flyway.table= #
flyway.target= #
flyway.url= # JDBC url of the database to migrate. If not set, the primary configured data source is
used.
flyway.user= # Login user of the database to migrate.
flyway.validate-on-migrate= #

# LIQUIBASE (LiquibaseProperties)
liquibase.change-log=classpath:/db/changelog/db.changelog-master.yaml # Change log configuration path.
liquibase.check-change-log-location=true # Check the change log location exists.
liquibase.contexts= # Comma-separated list of runtime contexts to use.
liquibase.default-schema= # Default database schema.
liquibase.drop-first=false # Drop the database schema first.
liquibase.enabled=true # Enable liquibase support.
liquibase.labels= # Comma-separated list of runtime labels to use.
liquibase.parameters.*= # Change log parameters.
liquibase.password= # Login password of the database to migrate.
liquibase.rollback-file= # File to which rollback SQL will be written when an update is performed.
liquibase.url= # JDBC url of the database to migrate. If not set, the primary configured data source is
used.
liquibase.user= # Login user of the database to migrate.

# COUCHBASE (CouchbaseProperties)
spring.couchbase.bootstrap-hosts= # Couchbase nodes (host or IP address) to bootstrap from.
spring.couchbase.bucket.name=default # Name of the bucket to connect to.
spring.couchbase.bucket.password= # Password of the bucket.
spring.couchbase.env.endpoints.key-value=1 # Number of sockets per node against the Key/value service.
spring.couchbase.env.endpoints.query=1 # Number of sockets per node against the Query (N1QL) service.
spring.couchbase.env.endpoints.view=1 # Number of sockets per node against the view service.
spring.couchbase.env.ssl.enabled= # Enable SSL support. Enabled automatically if a "keyStore" is
provided unless specified otherwise.
spring.couchbase.env.ssl.key-store= # Path to the JVM key store that holds the certificates.
spring.couchbase.env.ssl.key-store-password= # Password used to access the key store.
spring.couchbase.env.timeouts.connect=5000 # Bucket connections timeout in milliseconds.
spring.couchbase.env.timeouts.key-value=2500 # Blocking operations performed on a specific key timeout
in milliseconds.
spring.couchbase.env.timeouts.query=7500 # N1QL query operations timeout in milliseconds.
spring.couchbase.env.timeouts.socket-connect=1000 # Socket connect connections timeout in milliseconds.
spring.couchbase.env.timeouts.view=7500 # Regular and geospatial view operations timeout in
milliseconds.

# DAO (PersistenceExceptionTranslationAutoConfiguration)

```

```

spring.dao.exceptiontranslation.enabled=true # Enable the PersistenceExceptionTranslationPostProcessor.

# CASSANDRA (CassandraProperties)
spring.data.cassandra.cluster-name= # Name of the Cassandra cluster.
spring.data.cassandra.compression=none # Compression supported by the Cassandra binary protocol.
spring.data.cassandra.connect-timeout-millis= # Socket option: connection time out.
spring.data.cassandra.consistency-level= # Queries consistency level.
spring.data.cassandra.contact-points=localhost # Comma-separated list of cluster node addresses.
spring.data.cassandra.fetch-size= # Queries default fetch size.
spring.data.cassandra.keyspace-name= # Keyspace name to use.
spring.data.cassandra.load-balancing-policy= # Class name of the load balancing policy.
spring.data.cassandra.port= # Port of the Cassandra server.
spring.data.cassandra.password= # Login password of the server.
spring.data.cassandra.pool.heartbeat-interval=30 # Heartbeat interval (in seconds) after which a message
  is sent on an idle connection to make sure it's still alive.
spring.data.cassandra.pool.idle-timeout=120 # Idle timeout (in seconds) before an idle connection is
  removed.
spring.data.cassandra.pool.max-queue-size=256 # Maximum number of requests that get enqueued if no
  connection is available.
spring.data.cassandra.pool.pool-timeout=5000 # Pool timeout (in milliseconds) when trying to acquire a
  connection from a host's pool.
spring.data.cassandra.reactive-repositories.enabled=true # Enable Cassandra reactive repositories.
spring.data.cassandra.read-timeout-millis= # Socket option: read time out.
spring.data.cassandra.reconnection-policy= # Reconnection policy class.
spring.data.cassandra.repositories.enabled= # Enable Cassandra repositories.
spring.data.cassandra.retry-policy= # Class name of the retry policy.
spring.data.cassandra.serial-consistency-level= # Queries serial consistency level.
spring.data.cassandra.schema-action=none # Schema action to take at startup.
spring.data.cassandra.ssl=false # Enable SSL support.
spring.data.cassandra.username= # Login user of the server.

# DATA COUCHBASE (CouchbaseDataProperties)
spring.data.couchbase.auto-index=false # Automatically create views and indexes.
spring.data.couchbase.consistency=read-your-own-writes # Consistency to apply by default on generated
  queries.
spring.data.couchbase.repositories.enabled=true # Enable Couchbase repositories.

# ELASTICSEARCH (ElasticsearchProperties)
spring.data.elasticsearch.cluster-name=elasticsearch # Elasticsearch cluster name.
spring.data.elasticsearch.cluster-nodes= # Comma-separated list of cluster node addresses.
spring.data.elasticsearch.properties.*= # Additional properties used to configure the client.
spring.data.elasticsearch.repositories.enabled=true # Enable Elasticsearch repositories.

# DATA LDAP
spring.data.ldap.repositories.enabled=true # Enable LDAP repositories.

# MONGODB (MongoProperties)
spring.data.mongodb.authentication-database= # Authentication database name.
spring.data.mongodb.database=test # Database name.
spring.data.mongodb.field-naming-strategy= # Fully qualified name of the FieldNamingStrategy to use.
spring.data.mongodb.grid-fs-database= # GridFS database name.
spring.data.mongodb.host=localhost # Mongo server host. Cannot be set with uri.
spring.data.mongodb.password= # Login password of the mongo server. Cannot be set with uri.
spring.data.mongodb.port=27017 # Mongo server port. Cannot be set with uri.
spring.data.mongodb.reactive-repositories.enabled=true # Enable Mongo reactive repositories.
spring.data.mongodb.repositories.enabled=true # Enable Mongo repositories.
spring.data.mongodb.uri=mongodb://localhost/test # Mongo database URI. Cannot be set with host, port and
  credentials.
spring.data.mongodb.username= # Login user of the mongo server. Cannot be set with uri.

# DATA REDIS
spring.data.redis.repositories.enabled=true # Enable Redis repositories.

# NEO4J (Neo4jProperties)
spring.data.neo4j.auto-index=none # Auto index mode.
spring.data.neo4j.embedded.enabled=true # Enable embedded mode if the embedded driver is available.
spring.data.neo4j.open-in-view=false # Register OpenSessionInViewInterceptor. Binds a Neo4j Session to
  the thread for the entire processing of the request.
spring.data.neo4j.password= # Login password of the server.
spring.data.neo4j.repositories.enabled=true # Enable Neo4j repositories.
spring.data.neo4j.uri= # URI used by the driver. Auto-detected by default.

```

```

spring.data.neo4j.username= # Login user of the server.

# DATA REST (RepositoryRestProperties)
spring.data.rest.base-path= # Base path to be used by Spring Data REST to expose repository resources.
spring.data.rest.default-page-size= # Default size of pages.
spring.data.rest.detection-strategy=default # Strategy to use to determine which repositories get
exposed.
spring.data.rest.enable-enum-translation= # Enable enum value translation via the Spring Data REST
default resource bundle.
spring.data.rest.limit-param-name= # Name of the URL query string parameter that indicates how many
results to return at once.
spring.data.rest.max-page-size= # Maximum size of pages.
spring.data.rest.page-param-name= # Name of the URL query string parameter that indicates what page to
return.
spring.data.rest.return-body-on-create= # Return a response body after creating an entity.
spring.data.rest.return-body-on-update= # Return a response body after updating an entity.
spring.data.rest.sort-param-name= # Name of the URL query string parameter that indicates what direction
to sort results.

# SOLR (SolrProperties)
spring.data.solr.host=http://127.0.0.1:8983/solr # Solr host. Ignored if "zk-host" is set.
spring.data.solr.repositories.enabled=true # Enable Solr repositories.
spring.data.solr.zk-host= # ZooKeeper host address in the form HOST:PORT.

# DATA WEB (SpringDataWebProperties)
spring.data.web.pageable.default-page-size=20 # Default page size.
spring.data.web.pageable.page-parameter=page # Page index parameter name.
spring.data.web.pageable.size-parameter=size # Page size parameter name.
spring.data.web.sort.sort-parameter=sort # Sort parameter name.

# DATASOURCE (DataSourceAutoConfiguration & DataSourceProperties)
spring.datasource.continue-on-error=false # Do not stop if an error occurs while initializing the
database.
spring.datasource.data= # Data (DML) script resource references.
spring.datasource.data-username= # User of the database to execute DML scripts (if different).
spring.datasource.data-password= # Password of the database to execute DML scripts (if different).
spring.datasource.dbcp2.*= # Commons DBCP2 specific settings
spring.datasource.driver-class-name= # Fully qualified name of the JDBC driver. Auto-detected based on
the URL by default.
spring.datasource.generate-unique-name=false # Generate a random datasource name.
spring.datasource.hikari.*= # Hikari specific settings
spring.datasource.initialize=true # Populate the database using 'data.sql'.
spring.datasource.jmx-enabled=false # Enable JMX support (if provided by the underlying pool).
spring.datasource.jndi-name= # JNDI location of the datasource. Class, url, username & password are
ignored when set.
spring.datasource.name=testdb # Name of the datasource.
spring.datasource.password= # Login password of the database.
spring.datasource.platform=all # Platform to use in the schema resource (schema- $\{platform\}$ .sql).
spring.datasource.schema= # Schema (DDL) script resource references.
spring.datasource.schema-username= # User of the database to execute DDL scripts (if different).
spring.datasource.schema-password= # Password of the database to execute DDL scripts (if different).
spring.datasource.separator=; # Statement separator in SQL initialization scripts.
spring.datasource.sql-script-encoding= # SQL scripts encoding.
spring.datasource.tomcat.*= # Tomcat datasource specific settings
spring.datasource.type= # Fully qualified name of the connection pool implementation to use. By default,
it is auto-detected from the classpath.
spring.datasource.url= # JDBC url of the database.
spring.datasource.username= # Login user of the database.
spring.datasource.xa.data-source-class-name= # XA datasource fully qualified name.
spring.datasource.xa.properties= # Properties to pass to the XA data source.

# JEST (Elasticsearch HTTP client) (JestProperties)
spring.elasticsearch.jest.connection-timeout=3000 # Connection timeout in milliseconds.
spring.elasticsearch.jest.multi-threaded=true # Enable connection requests from multiple execution
threads.
spring.elasticsearch.jest.password= # Login password.
spring.elasticsearch.jest.proxy.host= # Proxy host the HTTP client should use.
spring.elasticsearch.jest.proxy.port= # Proxy port the HTTP client should use.
spring.elasticsearch.jest.read-timeout=3000 # Read timeout in milliseconds.
spring.elasticsearch.jest.uris=http://localhost:9200 # Comma-separated list of the Elasticsearch
instances to use.

```

```

spring.elasticsearch.jest.username= # Login user.

# H2 Web Console (H2ConsoleProperties)
spring.h2.console.enabled=false # Enable the console.
spring.h2.console.path=/h2-console # Path at which the console will be available.
spring.h2.console.settings.trace=false # Enable trace output.
spring.h2.console.settings.web-allow-others=false # Enable remote access.

# InfluxDB (InfluxDbProperties)
spring.influx.password= # Login password.
spring.influx.url= # Url of the InfluxDB instance to connect to.
spring.influx.user= # Login user.

# JOOQ (JooqAutoConfiguration)
spring.jooq.sql-dialect= # Sql dialect to use, auto-detected by default.

# JDBC (JdbcProperties)
spring.jdbc.template.fetch-size=-1 # Number of rows that should be fetched from the database when more
rows are needed.
spring.jdbc.template.max-rows=-1 # Maximum number of rows.
spring.jdbc.template.query-timeout=-1 # Query timeout in seconds.

# JPA (JpaBaseConfiguration, HibernateJpaAutoConfiguration)
spring.data.jpa.repositories.enabled=true # Enable JPA repositories.
spring.jpa.database= # Target database to operate on, auto-detected by default. Can be alternatively set
using the "databasePlatform" property.
spring.jpa.database-platform= # Name of the target database to operate on, auto-detected by default. Can
be alternatively set using the "Database" enum.
spring.jpa.generate-ddl=false # Initialize the schema on startup.
spring.jpa.hibernate.ddl-auto= # DDL mode. This is actually a shortcut for the "hibernate.hbm2ddl.auto"
property. Default to "create-drop" when using an embedded database, "none" otherwise.
spring.jpa.hibernate.naming.implicit-strategy= # Hibernate 5 implicit naming strategy fully qualified
name.
spring.jpa.hibernate.naming.physical-strategy= # Hibernate 5 physical naming strategy fully qualified
name.
spring.jpa.hibernate.use-new-id-generator-mappings= # Use Hibernate's newer IdentifierGenerator for
AUTO, TABLE and SEQUENCE.
spring.jpa.open-in-view=true # Register OpenEntityManagerInViewInterceptor. Binds a JPA EntityManager to
the thread for the entire processing of the request.
spring.jpa.properties.*= # Additional native properties to set on the JPA provider.
spring.jpa.show-sql=false # Enable logging of SQL statements.

# JTA (JtaAutoConfiguration)
spring.jta.enabled=true # Enable JTA support.
spring.jta.log-dir= # Transaction logs directory.
spring.jta.transaction-manager-id= # Transaction manager unique identifier.

# ATOMIKOS (AtomikosProperties)
spring.jta.atomikos.connectionfactory.borrow-connection-timeout=30 # Timeout, in seconds, for borrowing
connections from the pool.
spring.jta.atomikos.connectionfactory.ignore-session-transacted-flag=true # Whether or not to ignore the
transacted flag when creating session.
spring.jta.atomikos.connectionfactory.local-transaction-mode=false # Whether or not local transactions
are desired.
spring.jta.atomikos.connectionfactory.maintenance-interval=60 # The time, in seconds, between runs of
the pool's maintenance thread.
spring.jta.atomikos.connectionfactory.max-idle-time=60 # The time, in seconds, after which connections
are cleaned up from the pool.
spring.jta.atomikos.connectionfactory.max-lifetime=0 # The time, in seconds, that a connection can be
pooled for before being destroyed. 0 denotes no limit.
spring.jta.atomikos.connectionfactory.max-pool-size=1 # The maximum size of the pool.
spring.jta.atomikos.connectionfactory.min-pool-size=1 # The minimum size of the pool.
spring.jta.atomikos.connectionfactory.reap-timeout=0 # The reap timeout, in seconds, for borrowed
connections. 0 denotes no limit.
spring.jta.atomikos.connectionfactory.unique-resource-name=jmsConnectionFactory # The unique name used
to identify the resource during recovery.
spring.jta.atomikos.datasource.borrow-connection-timeout=30 # Timeout, in seconds, for borrowing
connections from the pool.
spring.jta.atomikos.datasource.default-isolation-level= # Default isolation level of connections
provided by the pool.

```

```

spring.jta.atomikos.datasource.login-timeout= # Timeout, in seconds, for establishing a database
connection.
spring.jta.atomikos.datasource.maintenance-interval=60 # The time, in seconds, between runs of the
pool's maintenance thread.
spring.jta.atomikos.datasource.max-idle-time=60 # The time, in seconds, after which connections are
cleaned up from the pool.
spring.jta.atomikos.datasource.max-lifetime=0 # The time, in seconds, that a connection can be pooled
for before being destroyed. 0 denotes no limit.
spring.jta.atomikos.datasource.max-pool-size=1 # The maximum size of the pool.
spring.jta.atomikos.datasource.min-pool-size=1 # The minimum size of the pool.
spring.jta.atomikos.datasource.reap-timeout=0 # The reap timeout, in seconds, for borrowed connections.
0 denotes no limit.
spring.jta.atomikos.datasource.test-query= # SQL query or statement used to validate a connection before
returning it.
spring.jta.atomikos.datasource.unique-resource-name=dataSource # The unique name used to identify the
resource during recovery.
spring.jta.atomikos.properties.allow-sub-transactions=true # Specify if sub-transactions are allowed.
spring.jta.atomikos.properties.checkpoint-interval=500 # Interval between checkpoints.
spring.jta.atomikos.properties.default-jta-timeout=10000 # Default timeout for JTA transactions.
spring.jta.atomikos.properties.enable-logging=true # Enable disk logging.
spring.jta.atomikos.properties.force-shutdown-on-vm-exit=false # Specify if a VM shutdown should trigger
forced shutdown of the transaction core.
spring.jta.atomikos.properties.log-base-dir= # Directory in which the log files should be stored.
spring.jta.atomikos.properties.log-base-name=tmlog # Transactions log file base name.
spring.jta.atomikos.properties.max-actives=50 # Maximum number of active transactions.
spring.jta.atomikos.properties.max-timeout=300000 # Maximum timeout (in milliseconds) that can be
allowed for transactions.
spring.jta.atomikos.properties.recovery.delay=10000 # Delay between two recovery scans.
spring.jta.atomikos.properties.recovery.forget-orphaned-log-entries-delay=86400000 # Delay after which
recovery can cleanup pending ('orphaned') log entries.
spring.jta.atomikos.properties.recovery.max-retries=5 # Number of retry attempts to commit the
transaction before throwing an exception.
spring.jta.atomikos.properties.recovery.retry-interval=10000 # Delay between retry attempts.
spring.jta.atomikos.properties.serial-jta-transactions=true # Specify if sub-transactions should be
joined when possible.
spring.jta.atomikos.properties.service= # Transaction manager implementation that should be started.
spring.jta.atomikos.properties.threaded-two-phase-commit=false # Use different (and concurrent) threads
for two-phase commit on the participating resources.
spring.jta.atomikos.properties.transaction-manager-unique-name= # Transaction manager's unique name.

# BITRONIX
spring.jta.bitronix.connectionfactory.acquire-increment=1 # Number of connections to create when growing
the pool.
spring.jta.bitronix.connectionfactory.acquisition-interval=1 # Time, in seconds, to wait before trying
to acquire a connection again after an invalid connection was acquired.
spring.jta.bitronix.connectionfactory.acquisition-timeout=30 # Timeout, in seconds, for acquiring
connections from the pool.
spring.jta.bitronix.connectionfactory.allow-local-transactions=true # Whether or not the transaction
manager should allow mixing XA and non-XA transactions.
spring.jta.bitronix.connectionfactory.apply-transaction-timeout=false # Whether or not the transaction
timeout should be set on the XAResource when it is enlisted.
spring.jta.bitronix.connectionfactory.automatic-enlisting-enabled=true # Whether or not resources should
be enlisted and delisted automatically.
spring.jta.bitronix.connectionfactory.cache-producers-consumers=true # Whether or not produces and
consumers should be cached.
spring.jta.bitronix.connectionfactory.defer-connection-release=true # Whether or not the provider can
run many transactions on the same connection and supports transaction interleaving.
spring.jta.bitronix.connectionfactory.ignore-recovery-failures=false # Whether or not recovery failures
should be ignored.
spring.jta.bitronix.connectionfactory.max-idle-time=60 # The time, in seconds, after which connections
are cleaned up from the pool.
spring.jta.bitronix.connectionfactory.max-pool-size=10 # The maximum size of the pool. 0 denotes no
limit.
spring.jta.bitronix.connectionfactory.min-pool-size=0 # The minimum size of the pool.
spring.jta.bitronix.connectionfactory.password= # The password to use to connect to the JMS provider.
spring.jta.bitronix.connectionfactory.share-transaction-connections=false # Whether or not connections
in the ACCESSIBLE state can be shared within the context of a transaction.
spring.jta.bitronix.connectionfactory.test-connections=true # Whether or not connections should be
tested when acquired from the pool.

```

```

spring.jta.bitronix.connectionfactory.two-pc-ordering-position=1 # The position that this
resource should take during two-phase commit (always first is Integer.MIN_VALUE, always last is
Integer.MAX_VALUE).
spring.jta.bitronix.connectionfactory.unique-name=jmsConnectionFactory # The unique name used to
identify the resource during recovery.
spring.jta.bitronix.connectionfactory.use-tm-join=true Whether or not TMJOIN should be used when
starting XAResources.
spring.jta.bitronix.connectionfactory.user= # The user to use to connect to the JMS provider.
spring.jta.bitronix.datasource.acquire-increment=1 # Number of connections to create when growing the
pool.
spring.jta.bitronix.datasource.acquisition-interval=1 # Time, in seconds, to wait before trying to
acquire a connection again after an invalid connection was acquired.
spring.jta.bitronix.datasource.acquisition-timeout=30 # Timeout, in seconds, for acquiring connections
from the pool.
spring.jta.bitronix.datasource.allow-local-transactions=true # Whether or not the transaction manager
should allow mixing XA and non-XA transactions.
spring.jta.bitronix.datasource.apply-transaction-timeout=false # Whether or not the transaction timeout
should be set on the XAResource when it is enlisted.
spring.jta.bitronix.datasource.automatic-enlisting-enabled=true # Whether or not resources should be
enlisted and delisted automatically.
spring.jta.bitronix.datasource.cursor-holdability= # The default cursor holdability for connections.
spring.jta.bitronix.datasource.defer-connection-release=true # Whether or not the database can run many
transactions on the same connection and supports transaction interleaving.
spring.jta.bitronix.datasource.enable-jdbc4-connection-test= # Whether or not Connection.isValid() is
called when acquiring a connection from the pool.
spring.jta.bitronix.datasource.ignore-recovery-failures=false # Whether or not recovery failures should
be ignored.
spring.jta.bitronix.datasource.isolation-level= # The default isolation level for connections.
spring.jta.bitronix.datasource.local-auto-commit= # The default auto-commit mode for local transactions.
spring.jta.bitronix.datasource.login-timeout= # Timeout, in seconds, for establishing a database
connection.
spring.jta.bitronix.datasource.max-idle-time=60 # The time, in seconds, after which connections are
cleaned up from the pool.
spring.jta.bitronix.datasource.max-pool-size=10 # The maximum size of the pool. 0 denotes no limit.
spring.jta.bitronix.datasource.min-pool-size=0 # The minimum size of the pool.
spring.jta.bitronix.datasource.prepared-statement-cache-size=0 # The target size of the prepared
statement cache. 0 disables the cache.
spring.jta.bitronix.datasource.share-transaction-connections=false # Whether or not connections in the
ACCESSIBLE state can be shared within the context of a transaction.
spring.jta.bitronix.datasource.test-query= # SQL query or statement used to validate a connection before
returning it.
spring.jta.bitronix.datasource.two-pc-ordering-position=1 # The position that this resource should take
during two-phase commit (always first is Integer.MIN_VALUE, always last is Integer.MAX_VALUE).
spring.jta.bitronix.datasource.unique-name=dataSource # The unique name used to identify the resource
during recovery.
spring.jta.bitronix.datasource.use-tm-join=true Whether or not TMJOIN should be used when starting
XAResources.
spring.jta.bitronix.properties.allow-multiple-lrc=false # Allow multiple LRC resources to be enlisted
into the same transaction.
spring.jta.bitronix.properties.asynchronous2-pc=false # Enable asynchronously execution of two phase
commit.
spring.jta.bitronix.properties.background-recovery-interval-seconds=60 # Interval in seconds at which to
run the recovery process in the background.
spring.jta.bitronix.properties.current-node-only-recovery=true # Recover only the current node.
spring.jta.bitronix.properties.debug-zero-resource-transaction=false # Log the creation and commit call
stacks of transactions executed without a single enlisted resource.
spring.jta.bitronix.properties.default-transaction-timeout=60 # Default transaction timeout in seconds.
spring.jta.bitronix.properties.disable-jmx=false # Enable JMX support.
spring.jta.bitronix.properties.exception-analyzer= # Set the fully qualified name of the exception
analyzer implementation to use.
spring.jta.bitronix.properties.filter-log-status=false # Enable filtering of logs so that only mandatory
logs are written.
spring.jta.bitronix.properties.force-batching-enabled=true # Set if disk forces are batched.
spring.jta.bitronix.properties.forced-write-enabled=true # Set if logs are forced to disk.
spring.jta.bitronix.properties.graceful-shutdown-interval=60 # Maximum amount of seconds the TM will
wait for transactions to get done before aborting them at shutdown time.
spring.jta.bitronix.properties.jndi-transaction-synchronization-registry-name= # JNDI name of the
TransactionSynchronizationRegistry.
spring.jta.bitronix.properties.jndi-user-transaction-name= # JNDI name of the UserTransaction.
spring.jta.bitronix.properties.journal=disk # Name of the journal. Can be 'disk', 'null' or a class
name.

```

```

spring.jta.bitronix.properties.log-part1-filename=btm1.tlog # Name of the first fragment of the journal.
spring.jta.bitronix.properties.log-part2-filename=btm2.tlog # Name of the second fragment of the
journal.
spring.jta.bitronix.properties.max-log-size-in-mb=2 # Maximum size in megabytes of the journal
fragments.
spring.jta.bitronix.properties.resource-configuration-filename= # ResourceLoader configuration file
name.
spring.jta.bitronix.properties.server-id= # ASCII ID that must uniquely identify this TM instance.
Default to the machine's IP address.
spring.jta.bitronix.properties.skip-corrupted-logs=false # Skip corrupted transactions log entries.
spring.jta.bitronix.properties.warn-about-zero-resource-transaction=true # Log a warning for
transactions executed without a single enlisted resource.

# NARAYANA (NarayanaProperties)
spring.jta.narayana.default-timeout=60 # Transaction timeout in seconds.
spring.jta.narayana.expiry-
scanners=com.arjuna.ats.internal.arjuna.recovery.ExpiredTransactionStatusManagerScanner # Comma-
separated list of expiry scanners.
spring.jta.narayana.log-dir= # Transaction object store directory.
spring.jta.narayana.one-phase-commit=true # Enable one phase commit optimisation.
spring.jta.narayana.periodic-recovery-period=120 # Interval in which periodic recovery scans are
performed in seconds.
spring.jta.narayana.recovery-backoff-period=10 # Back off period between first and second phases of the
recovery scan in seconds.
spring.jta.narayana.recovery-db-pass= # Database password to be used by recovery manager.
spring.jta.narayana.recovery-db-user= # Database username to be used by recovery manager.
spring.jta.narayana.recovery-jms-pass= # JMS password to be used by recovery manager.
spring.jta.narayana.recovery-jms-user= # JMS username to be used by recovery manager.
spring.jta.narayana.recovery-modules= # Comma-separated list of recovery modules.
spring.jta.narayana.transaction-manager-id=1 # Unique transaction manager id.
spring.jta.narayana.xa-resource-orphan-filters= # Comma-separated list of orphan filters.

# EMBEDDED MONGODB (EmbeddedMongoProperties)
spring.mongodb.embedded.features=SYNC_DELAY # Comma-separated list of features to enable.
spring.mongodb.embedded.storage.database-dir= # Directory used for data storage.
spring.mongodb.embedded.storage.oplog-size= # Maximum size of the oplog in megabytes.
spring.mongodb.embedded.storage.repl-set-name= # Name of the replica set.
spring.mongodb.embedded.version=2.6.10 # Version of Mongo to use.

# REDIS (RedisProperties)
spring.redis.cluster.max-redirects= # Maximum number of redirects to follow when executing commands
across the cluster.
spring.redis.cluster.nodes= # Comma-separated list of "host:port" pairs to bootstrap from.
spring.redis.database=0 # Database index used by the connection factory.
spring.redis.url= # Connection URL, will override host, port and password (user will be ignored), e.g.
redis://user:password@example.com:6379
spring.redis.host=localhost # Redis server host.
spring.redis.jedis.pool.max-active=8 # Max number of connections that can be allocated by the pool at a
given time. Use a negative value for no limit.
spring.redis.jedis.pool.max-idle=8 # Max number of "idle" connections in the pool. Use a negative value
to indicate an unlimited number of idle connections.
spring.redis.jedis.pool.max-wait=-1 # Maximum amount of time (in milliseconds) a connection allocation
should block before throwing an exception when the pool is exhausted. Use a negative value to block
indefinitely.
spring.redis.jedis.pool.min-idle=0 # Target for the minimum number of idle connections to maintain in
the pool. This setting only has an effect if it is positive.
spring.redis.lettuce.pool.max-active=8 # Max number of connections that can be allocated by the pool at
a given time. Use a negative value for no limit.
spring.redis.lettuce.pool.max-idle=8 # Max number of "idle" connections in the pool. Use a negative
value to indicate an unlimited number of idle connections.
spring.redis.lettuce.pool.max-wait=-1 # Maximum amount of time (in milliseconds) a connection allocation
should block before throwing an exception when the pool is exhausted. Use a negative value to block
indefinitely.
spring.redis.lettuce.pool.min-idle=0 # Target for the minimum number of idle connections to maintain in
the pool. This setting only has an effect if it is positive.
spring.redis.lettuce.shutdown-timeout=100 # Shutdown timeout in milliseconds.
spring.redis.password= # Login password of the redis server.
spring.redis.port=6379 # Redis server port.
spring.redis.sentinel.master= # Name of Redis server.
spring.redis.sentinel.nodes= # Comma-separated list of host:port pairs.
spring.redis.ssl=false # Enable SSL support.

```



```

spring.redis.timeout=0 # Connection timeout in milliseconds.

# TRANSACTION (TransactionProperties)
spring.transaction.default-timeout= # Default transaction timeout in seconds.
spring.transaction.rollback-on-commit-failure= # Perform the rollback on commit failures.

# -----
# INTEGRATION PROPERTIES
# -----

# ACTIVEMQ (ActiveMQProperties)
spring.activemq.broker-url= # URL of the ActiveMQ broker. Auto-generated by default.
spring.activemq.close-timeout=15000 # Time to wait, in milliseconds, before considering a close
complete.
spring.activemq.in-memory=true # Specify if the default broker URL should be in memory. Ignored if an
explicit broker has been specified.
spring.activemq.non-blocking-redelivery=false # Do not stop message delivery before re-delivering
messages from a rolled back transaction. This implies that message order will not be preserved when
this is enabled.
spring.activemq.password= # Login password of the broker.
spring.activemq.send-timeout=0 # Time to wait, in milliseconds, on Message sends for a response. Set it
to 0 to indicate to wait forever.
spring.activemq.user= # Login user of the broker.
spring.activemq.packages.trust-all= # Trust all packages.
spring.activemq.packages.trusted= # Comma-separated list of specific packages to trust (when not
trusting all packages).
spring.activemq.pool.block-if-full=true # Block when a connection is requested and the pool is full. Set
it to false to throw a "JMSEException" instead.
spring.activemq.pool.block-if-full-timeout=-1 # Blocking period, in milliseconds, before throwing an
exception if the pool is still full.
spring.activemq.pool.create-connection-on-startup=true # Create a connection on startup. Can be used to
warm-up the pool on startup.
spring.activemq.pool.enabled=false # Whether a PooledConnectionFactory should be created instead of a
regular ConnectionFactory.
spring.activemq.pool.expiry-timeout=0 # Connection expiration timeout in milliseconds.
spring.activemq.pool.idle-timeout=30000 # Connection idle timeout in milliseconds.
spring.activemq.pool.max-connections=1 # Maximum number of pooled connections.
spring.activemq.pool.maximum-active-session-per-connection=500 # Maximum number of active sessions per
connection.
spring.activemq.pool.reconnect-on-exception=true # Reset the connection when a "JMSEException" occurs.
spring.activemq.pool.time-between-expiration-check=-1 # Time to sleep, in milliseconds, between runs of
the idle connection eviction thread. When negative, no idle connection eviction thread runs.
spring.activemq.pool.use-anonymous-producers=true # Use only one anonymous "MessageProducer" instance.
Set it to false to create one "MessageProducer" every time one is required.

# ARTEMIS (ArtemisProperties)
spring.artemis.embedded.cluster-password= # Cluster password. Randomly generated on startup by default.
spring.artemis.embedded.data-directory= # Journal file directory. Not necessary if persistence is turned
off.
spring.artemis.embedded.enabled=true # Enable embedded mode if the Artemis server APIs are available.
spring.artemis.embedded.persistent=false # Enable persistent store.
spring.artemis.embedded.queues= # Comma-separated list of queues to create on startup.
spring.artemis.embedded.server-id= # Server id. By default, an auto-incremented counter is used.
spring.artemis.embedded.topics= # Comma-separated list of topics to create on startup.
spring.artemis.host=localhost # Artemis broker host.
spring.artemis.mode= # Artemis deployment mode, auto-detected by default.
spring.artemis.password= # Login password of the broker.
spring.artemis.port=61616 # Artemis broker port.
spring.artemis.user= # Login user of the broker.

# SPRING BATCH (BatchProperties)
spring.batch.initializer.enabled= # Create the required batch tables on startup if necessary. Enabled
automatically if no custom table prefix is set or if a custom schema is configured.
spring.batch.job.enabled=true # Execute all Spring Batch jobs in the context on startup.
spring.batch.job.names= # Comma-separated list of job names to execute on startup (For instance
`job1,job2`). By default, all Jobs found in the context are executed.
spring.batch.schema=classpath:org/springframework/batch/core/schema-@@platform@@.sql # Path to the SQL
file to use to initialize the database schema.
spring.batch.table-prefix= # Table prefix for all the batch meta-data tables.

```

```

# SPRING INTEGRATION (IntegrationProperties)
spring.integration.jdbc.initializer.enabled=false # Create the required integration tables on startup.
spring.integration.jdbc.schema=classpath:org/springframework/integration/jdbc/schema-@@platform@@.sql #
Path to the SQL file to use to initialize the database schema.

# JMS (JmsProperties)
spring.jms.jndi-name= # Connection factory JNDI name. When set, takes precedence to others connection
factory auto-configurations.
spring.jms.listener.acknowledge-mode= # Acknowledge mode of the container. By default, the listener is
transacted with automatic acknowledgment.
spring.jms.listener.auto-startup=true # Start the container automatically on startup.
spring.jms.listener.concurrency= # Minimum number of concurrent consumers.
spring.jms.listener.max-concurrency= # Maximum number of concurrent consumers.
spring.jms.pub-sub-domain=false # Specify if the default destination type is topic.
spring.jms.template.default-destination= # Default destination to use on send/receive operations that do
not have a destination parameter.
spring.jms.template.delivery-delay= # Delivery delay to use for send calls in milliseconds.
spring.jms.template.delivery-mode= # Delivery mode. Enable QoS when set.
spring.jms.template.priority= # Priority of a message when sending. Enable QoS when set.
spring.jms.template.qos-enabled= # Enable explicit QoS when sending a message.
spring.jms.template.receive-timeout= # Timeout to use for receive calls in milliseconds.
spring.jms.template.time-to-live= # Time-to-live of a message when sending in milliseconds. Enable QoS
when set.

# APACHE KAFKA (KafkaProperties)
spring.kafka.bootstrap-servers= # Comma-delimited list of host:port pairs to use for establishing the
initial connection to the Kafka cluster.
spring.kafka.client-id= # Id to pass to the server when making requests; used for server-side logging.
spring.kafka.consumer.auto-commit-interval= # Frequency in milliseconds that the consumer offsets are
auto-committed to Kafka if 'enable.auto.commit' true.
spring.kafka.consumer.auto-offset-reset= # What to do when there is no initial offset in Kafka or if the
current offset does not exist any more on the server.
spring.kafka.consumer.bootstrap-servers= # Comma-delimited list of host:port pairs to use for
establishing the initial connection to the Kafka cluster.
spring.kafka.consumer.client-id= # Id to pass to the server when making requests; used for server-side
logging.
spring.kafka.consumer.enable-auto-commit= # If true the consumer's offset will be periodically committed
in the background.
spring.kafka.consumer.fetch-max-wait= # Maximum amount of time in milliseconds the server will
block before answering the fetch request if there isn't sufficient data to immediately satisfy the
requirement given by "fetch.min.bytes".
spring.kafka.consumer.fetch-min-size= # Minimum amount of data the server should return for a fetch
request in bytes.
spring.kafka.consumer.group-id= # Unique string that identifies the consumer group this consumer belongs
to.
spring.kafka.consumer.heartbeat-interval= # Expected time in milliseconds between heartbeats to the
consumer coordinator.
spring.kafka.consumer.key-deserializer= # Deserializer class for keys.
spring.kafka.consumer.max-poll-records= # Maximum number of records returned in a single call to poll().
spring.kafka.consumer.properties.*= # Additional consumer-specific properties used to configure the
client.
spring.kafka.consumer.ssl.key-password= # Password of the private key in the key store file.
spring.kafka.consumer.ssl.keystore-location= # Location of the key store file.
spring.kafka.consumer.ssl.keystore-password= # Store password for the key store file.
spring.kafka.consumer.ssl.truststore-location= # Location of the trust store file.
spring.kafka.consumer.ssl.truststore-password= # Store password for the trust store file.
spring.kafka.consumer.value-deserializer= # Deserializer class for values.
spring.kafka.jaas.control-flag=required # Control flag for login configuration.
spring.kafka.jaas.enabled= # Enable JAAS configuration.
spring.kafka.jaas.login-module=com.sun.security.auth.module.Krb5LoginModule # Login module.
spring.kafka.jaas.options= # Additional JAAS options.
spring.kafka.listener.ack-count= # Number of records between offset commits when ackMode is "COUNT" or
"COUNT_TIME".
spring.kafka.listener.ack-mode= # Listener AckMode; see the spring-kafka documentation.
spring.kafka.listener.ack-time= # Time in milliseconds between offset commits when ackMode is "TIME" or
"COUNT_TIME".
spring.kafka.listener.concurrency= # Number of threads to run in the listener containers.
spring.kafka.listener.poll-timeout= # Timeout in milliseconds to use when polling the consumer.
spring.kafka.listener.type=single # Listener type.

```

```

spring.kafka.producer.acks= # Number of acknowledgments the producer requires the leader to have
    received before considering a request complete.
spring.kafka.producer.batch-size= # Number of records to batch before sending.
spring.kafka.producer.bootstrap-servers= # Comma-delimited list of host:port pairs to use for
    establishing the initial connection to the Kafka cluster.
spring.kafka.producer.buffer-memory= # Total bytes of memory the producer can use to buffer records
    waiting to be sent to the server.
spring.kafka.producer.client-id= # Id to pass to the server when making requests; used for server-side
    logging.
spring.kafka.producer.compression-type= # Compression type for all data generated by the producer.
spring.kafka.producer.key-serializer= # Serializer class for keys.
spring.kafka.producer.properties.*= # Additional producer-specific properties used to configure the
    client.
spring.kafka.producer.retries= # When greater than zero, enables retrying of failed sends.
spring.kafka.producer.ssl.key-password= # Password of the private key in the key store file.
spring.kafka.producer.ssl.keystore-location= # Location of the key store file.
spring.kafka.producer.ssl.keystore-password= # Store password for the key store file.
spring.kafka.producer.ssl.truststore-location= # Location of the trust store file.
spring.kafka.producer.ssl.truststore-password= # Store password for the trust store file.
spring.kafka.producer.value-serializer= # Serializer class for values.
spring.kafka.properties.*= # Additional properties, common to producers and consumers, used to configure
    the client.
spring.kafka.ssl.key-password= # Password of the private key in the key store file.
spring.kafka.ssl.keystore-location= # Location of the key store file.
spring.kafka.ssl.keystore-password= # Store password for the key store file.
spring.kafka.ssl.truststore-location= # Location of the trust store file.
spring.kafka.ssl.truststore-password= # Store password for the trust store file.
spring.kafka.template.default-topic= # Default topic to which messages will be sent.

# RABBIT (RabbitProperties)
spring.rabbitmq.addresses= # Comma-separated list of addresses to which the client should connect.
spring.rabbitmq.cache.channel.checkout-timeout= # Number of milliseconds to wait to obtain a channel if
    the cache size has been reached.
spring.rabbitmq.cache.channel.size= # Number of channels to retain in the cache.
spring.rabbitmq.cache.connection.mode=channel # Connection factory cache mode.
spring.rabbitmq.cache.connection.size= # Number of connections to cache.
spring.rabbitmq.connection-timeout= # Connection timeout, in milliseconds; zero for infinite.
spring.rabbitmq.dynamic=true # Create an AmqpAdmin bean.
spring.rabbitmq.host=localhost # RabbitMQ host.
spring.rabbitmq.listener.direct.acknowledge-mode= # Acknowledge mode of container.
spring.rabbitmq.listener.direct.auto-startup=true # Start the container automatically on startup.
spring.rabbitmq.listener.direct.consumers-per-queue= # Number of consumers per queue.
spring.rabbitmq.listener.direct.default-requeue-rejected= # Whether rejected deliveries are requeued by
    default; default true.
spring.rabbitmq.listener.direct.idle-event-interval= # How often idle container events should be
    published in milliseconds.
spring.rabbitmq.listener.direct.prefetch= # Number of messages to be handled in a single request. It
    should be greater than or equal to the transaction size (if used).
spring.rabbitmq.listener.direct.retry.enabled=false # Whether or not publishing retries are enabled.
spring.rabbitmq.listener.direct.retry.initial-interval=1000 # Interval between the first and second
    attempt to publish or deliver a message.
spring.rabbitmq.listener.direct.retry.max-attempts=3 # Maximum number of attempts to publish or deliver
    a message.
spring.rabbitmq.listener.direct.retry.max-interval=10000 # Maximum interval between attempts.
spring.rabbitmq.listener.direct.retry.multiplier=1 # A multiplier to apply to the previous retry
    interval.
spring.rabbitmq.listener.direct.retry.stateless=true # Whether or not retries are stateless or stateful.
spring.rabbitmq.listener.simple.acknowledge-mode= # Acknowledge mode of container.
spring.rabbitmq.listener.simple.auto-startup=true # Start the container automatically on startup.
spring.rabbitmq.listener.simple.concurrency= # Minimum number of listener invoker threads.
spring.rabbitmq.listener.simple.default-requeue-rejected= # Whether or not to requeue delivery failures.
spring.rabbitmq.listener.simple.idle-event-interval= # How often idle container events should be
    published in milliseconds.
spring.rabbitmq.listener.simple.max-concurrency= # Maximum number of listener invoker.
spring.rabbitmq.listener.simple.prefetch= # Number of messages to be handled in a single request. It
    should be greater than or equal to the transaction size (if used).
spring.rabbitmq.listener.simple.retry.enabled=false # Whether or not publishing retries are enabled.
spring.rabbitmq.listener.simple.retry.initial-interval=1000 # Interval between the first and second
    attempt to deliver a message.
spring.rabbitmq.listener.simple.retry.max-attempts=3 # Maximum number of attempts to deliver a message.
spring.rabbitmq.listener.simple.retry.max-interval=10000 # Maximum interval between attempts.

```

```

spring.rabbitmq.listener.simple.retry.multiplier=1.0 # A multiplier to apply to the previous delivery
  retry interval.
spring.rabbitmq.listener.simple.retry.stateless=true # Whether or not retry is stateless or stateful.
spring.rabbitmq.listener.simple.transaction-size= # Number of messages to be processed in a transaction;
  number of messages between acks. For best results it should be less than or equal to the prefetch
  count.
spring.rabbitmq.listener.type=simple # Listener container type.
spring.rabbitmq.password= # Login to authenticate against the broker.
spring.rabbitmq.port=5672 # RabbitMQ port.
spring.rabbitmq.publisher-confirms=false # Enable publisher confirms.
spring.rabbitmq.publisher-returns=false # Enable publisher returns.
spring.rabbitmq.requested-heartbeat= # Requested heartbeat timeout, in seconds; zero for none.
spring.rabbitmq.ssl.enabled=false # Enable SSL support.
spring.rabbitmq.ssl.key-store= # Path to the key store that holds the SSL certificate.
spring.rabbitmq.ssl.key-store-password= # Password used to access the key store.
spring.rabbitmq.ssl.trust-store= # Trust store that holds SSL certificates.
spring.rabbitmq.ssl.trust-store-password= # Password used to access the trust store.
spring.rabbitmq.ssl.algorithm= # SSL algorithm to use. By default configure by the rabbit client
  library.
spring.rabbitmq.template.mandatory=false # Enable mandatory messages.
spring.rabbitmq.template.receive-timeout=0 # Timeout for `receive()` methods.
spring.rabbitmq.template.reply-timeout=5000 # Timeout for `sendAndReceive()` methods.
spring.rabbitmq.template.retry.enabled=false # Set to true to enable retries in the `RabbitTemplate`.
spring.rabbitmq.template.retry.initial-interval=1000 # Interval between the first and second attempt to
  publish a message.
spring.rabbitmq.template.retry.max-attempts=3 # Maximum number of attempts to publish a message.
spring.rabbitmq.template.retry.max-interval=10000 # Maximum number of attempts to publish a message.
spring.rabbitmq.template.retry.multiplier=1.0 # A multiplier to apply to the previous publishing retry
  interval.
spring.rabbitmq.username= # Login user to authenticate to the broker.
spring.rabbitmq.virtual-host= # Virtual host to use when connecting to the broker.

# -----
# ACTUATOR PROPERTIES
# -----

# ENDPOINTS (AbstractEndpoint subclasses)
endpoints.enabled=true # Enable endpoints.
endpoints.sensitive= # Default endpoint sensitive setting.
endpoints.actuator.enabled=true # Enable the endpoint.
endpoints.actuator.path= # Endpoint URL path.
endpoints.actuator.sensitive=false # Enable security on the endpoint.
endpoints.auditevents.enabled= # Enable the endpoint.
endpoints.auditevents.path= # Endpoint path.
endpoints.auditevents.sensitive=false # Enable security on the endpoint.
endpoints.autoconfig.enabled= # Enable the endpoint.
endpoints.autoconfig.id= # Endpoint identifier.
endpoints.autoconfig.path= # Endpoint path.
endpoints.autoconfig.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.beans.enabled= # Enable the endpoint.
endpoints.beans.id= # Endpoint identifier.
endpoints.beans.path= # Endpoint path.
endpoints.beans.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.configprops.enabled= # Enable the endpoint.
endpoints.configprops.id= # Endpoint identifier.
endpoints.configprops.keys-to-sanitize=password,secret,key,token,.*credentials.*,vcap_services # Keys
  that should be sanitized. Keys can be simple strings that the property ends with or regex expressions.
endpoints.configprops.path= # Endpoint path.
endpoints.configprops.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.docs.curies.enabled=false # Enable the curie generation.
endpoints.docs.enabled=true # Enable actuator docs endpoint.
endpoints.docs.path=/docs #
endpoints.docs.sensitive=false #
endpoints.dump.enabled= # Enable the endpoint.
endpoints.dump.id= # Endpoint identifier.
endpoints.dump.path= # Endpoint path.
endpoints.dump.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.env.enabled= # Enable the endpoint.
endpoints.env.id= # Endpoint identifier.

```

```

endpoints.env.keys-to-sanitize=password,secret,key,token,.*credentials.*,vcap_services # Keys that
  should be sanitized. Keys can be simple strings that the property ends with or regex expressions.
endpoints.env.path= # Endpoint path.
endpoints.env.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.flyway.enabled= # Enable the endpoint.
endpoints.flyway.id= # Endpoint identifier.
endpoints.flyway.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.health.enabled= # Enable the endpoint.
endpoints.health.id= # Endpoint identifier.
endpoints.health.mapping.*= # Mapping of health statuses to HttpStatus codes. By default, registered
  health statuses map to sensible defaults (i.e. UP maps to 200).
endpoints.health.path= # Endpoint path.
endpoints.health.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.health.time-to-live=1000 # Time to live for cached result, in milliseconds.
endpoints.heapdump.enabled= # Enable the endpoint.
endpoints.heapdump.path= # Endpoint path.
endpoints.heapdump.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.hypermedia.enabled=false # Enable hypermedia support for endpoints.
endpoints.info.enabled= # Enable the endpoint.
endpoints.info.id= # Endpoint identifier.
endpoints.info.path= # Endpoint path.
endpoints.info.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.jolokia.enabled=true # Enable Jolokia endpoint.
endpoints.jolokia.path=/jolokia # Endpoint URL path.
endpoints.jolokia.sensitive=true # Enable security on the endpoint.
endpoints.liquibase.enabled= # Enable the endpoint.
endpoints.liquibase.id= # Endpoint identifier.
endpoints.liquibase.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.logfile.enabled=true # Enable the endpoint.
endpoints.logfile.external-file= # External Logfile to be accessed.
endpoints.logfile.path=/logfile # Endpoint URL path.
endpoints.logfile.sensitive=true # Enable security on the endpoint.
endpoints.loggers.enabled=true # Enable the endpoint.
endpoints.loggers.id= # Endpoint identifier.
endpoints.loggers.path=/logfile # Endpoint path.
endpoints.loggers.sensitive=true # Mark if the endpoint exposes sensitive information.
endpoints.mappings.enabled= # Enable the endpoint.
endpoints.mappings.id= # Endpoint identifier.
endpoints.mappings.path= # Endpoint path.
endpoints.mappings.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.metrics.enabled= # Enable the endpoint.
endpoints.metrics.filter.enabled=true # Enable the metrics servlet filter.
endpoints.metrics.filter.gauge-submissions=merged # Http filter gauge submissions (merged, per-http-
  method)
endpoints.metrics.filter.counter-submissions=merged # Http filter counter submissions (merged, per-http-
  method)
endpoints.metrics.id= # Endpoint identifier.
endpoints.metrics.path= # Endpoint path.
endpoints.metrics.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.shutdown.enabled= # Enable the endpoint.
endpoints.shutdown.id= # Endpoint identifier.
endpoints.shutdown.path= # Endpoint path.
endpoints.shutdown.sensitive= # Mark if the endpoint exposes sensitive information.
endpoints.trace.enabled= # Enable the endpoint.
endpoints.trace.filter.enabled=true # Enable the trace servlet filter.
endpoints.trace.id= # Endpoint identifier.
endpoints.trace.path= # Endpoint path.
endpoints.trace.sensitive= # Mark if the endpoint exposes sensitive information.

# ENDPOINTS CORS CONFIGURATION (EndpointCorsProperties)
endpoints.cors.allow-credentials= # Set whether credentials are supported. When not set, credentials are
  not supported.
endpoints.cors.allowed-headers= # Comma-separated list of headers to allow in a request. '*' allows all
  headers.
endpoints.cors.allowed-methods=GET # Comma-separated list of methods to allow. '*' allows all methods.
endpoints.cors.allowed-origins= # Comma-separated list of origins to allow. '*' allows all origins. When
  not set, CORS support is disabled.
endpoints.cors.exposed-headers= # Comma-separated list of headers to include in a response.
endpoints.cors.max-age=1800 # How long, in seconds, the response from a pre-flight request can be cached
  by clients.

```

```

# JMX ENDPOINT (EndpointMBeanExportProperties)
endpoints.jmx.domain= # JMX domain name. Initialized with the value of 'spring.jmx.default-domain' if
set.
endpoints.jmx.enabled=true # Enable JMX export of all endpoints.
endpoints.jmx.static-names= # Additional static properties to append to all ObjectNames of MBeans
representing Endpoints.
endpoints.jmx.unique-names=false # Ensure that ObjectNames are modified in case of conflict.

# JOLOKIA (JolokiaProperties)
jolokia.config.*= # See Jolokia manual

# MANAGEMENT HTTP SERVER (ManagementServerProperties)
management.add-application-context-header=false # Add the "X-Application-Context" HTTP header in each
response.
management.address= # Network address that the management endpoints should bind to.
management.context-path= # Management endpoint context-path. For instance `/actuator`
management.cloudfoundry.enabled= # Enable extended Cloud Foundry actuator endpoints
management.cloudfoundry.skip-ssl-validation= # Skip SSL verification for Cloud Foundry actuator endpoint
security calls
management.port= # Management endpoint HTTP port. Uses the same port as the application by default.
Configure a different port to use management-specific SSL.
management.security.enabled=true # Enable security.
management.security.roles=ACTUATOR # Comma-separated list of roles that can access the management
endpoint.
management.security.sessions=stateless # Session creating policy to use (always, never, if_required,
stateless).
management.ssl.ciphers= # Supported SSL ciphers. Requires a custom management.port.
management.ssl.client-auth= # Whether client authentication is wanted ("want") or needed ("need").
Requires a trust store. Requires a custom management.port.
management.ssl.enabled= # Enable SSL support. Requires a custom management.port.
management.ssl.enabled-protocols= # Enabled SSL protocols. Requires a custom management.port.
management.ssl.key-alias= # Alias that identifies the key in the key store. Requires a custom
management.port.
management.ssl.key-password= # Password used to access the key in the key store. Requires a custom
management.port.
management.ssl.key-store= # Path to the key store that holds the SSL certificate (typically a jks file).
Requires a custom management.port.
management.ssl.key-store-password= # Password used to access the key store. Requires a custom
management.port.
management.ssl.key-store-provider= # Provider for the key store. Requires a custom management.port.
management.ssl.key-store-type= # Type of the key store. Requires a custom management.port.
management.ssl.protocol=TLS # SSL protocol to use. Requires a custom management.port.
management.ssl.trust-store= # Trust store that holds SSL certificates. Requires a custom
management.port.
management.ssl.trust-store-password= # Password used to access the trust store. Requires a custom
management.port.
management.ssl.trust-store-provider= # Provider for the trust store. Requires a custom management.port.
management.ssl.trust-store-type= # Type of the trust store. Requires a custom management.port.

# HEALTH INDICATORS
management.health.db.enabled=true # Enable database health check.
management.health.cassandra.enabled=true # Enable cassandra health check.
management.health.couchbase.enabled=true # Enable couchbase health check.
management.health.defaults.enabled=true # Enable default health indicators.
management.health.diskspace.enabled=true # Enable disk space health check.
management.health.diskspace.path= # Path used to compute the available disk space.
management.health.diskspace.threshold=0 # Minimum disk space that should be available, in bytes.
management.health.elasticsearch.enabled=true # Enable elasticsearch health check.
management.health.elasticsearch.indices= # Comma-separated index names.
management.health.elasticsearch.response-timeout=100 # The time, in milliseconds, to wait for a response
from the cluster.
management.health.jms.enabled=true # Enable JMS health check.
management.health.ldap.enabled=true # Enable LDAP health check.
management.health.mail.enabled=true # Enable Mail health check.
management.health.mongo.enabled=true # Enable MongoDB health check.
management.health.neo4j.enabled=true # Enable Neo4j health check.
management.health.rabbit.enabled=true # Enable RabbitMQ health check.
management.health.redis.enabled=true # Enable Redis health check.
management.health.solr.enabled=true # Enable Solr health check.
management.health.status.order=DOWN, OUT_OF_SERVICE, UP, UNKNOWN # Comma-separated list of health
statuses in order of severity.

```

```

# INFO CONTRIBUTORS (InfoContributorProperties)
management.info.build.enabled=true # Enable build info.
management.info.defaults.enabled=true # Enable default info contributors.
management.info.env.enabled=true # Enable environment info.
management.info.git.enabled=true # Enable git info.
management.info.git.mode=simple # Mode to use to expose git information.

# TRACING (TraceProperties)
management.trace.include=request-headers,response-headers,cookies,errors # Items to be included in the
trace.

# METRICS EXPORT (MetricExportProperties)
spring.metrics.export.aggregate.key-pattern= # Pattern that tells the aggregator what to do with the
keys from the source repository.
spring.metrics.export.aggregate.prefix= # Prefix for global repository if active.
spring.metrics.export.delay-millis=5000 # Delay in milliseconds between export ticks. Metrics are
exported to external sources on a schedule with this delay.
spring.metrics.export.enabled=true # Flag to enable metric export (assuming a MetricWriter is
available).
spring.metrics.export.excludes= # List of patterns for metric names to exclude. Applied after the
includes.
spring.metrics.export.includes= # List of patterns for metric names to include.
spring.metrics.export.redis.key=keys.spring.metrics # Key for redis repository export (if active).
spring.metrics.export.redis.prefix=spring.metrics # Prefix for redis repository if active.
spring.metrics.export.send-latest= # Flag to switch off any available optimizations based on not
exporting unchanged metric values.
spring.metrics.export.statsd.host= # Host of a statsd server to receive exported metrics.
spring.metrics.export.statsd.port=8125 # Port of a statsd server to receive exported metrics.
spring.metrics.export.statsd.prefix= # Prefix for statsd exported metrics.
spring.metrics.export.triggers.*= # Specific trigger properties per MetricWriter bean name.

# -----
# DEVTOOLS PROPERTIES
# -----

# DEVTOOLS (DevToolsProperties)
spring.devtools.livereload.enabled=true # Enable a livereload.com compatible server.
spring.devtools.livereload.port=35729 # Server port.
spring.devtools.restart.additional-exclude= # Additional patterns that should be excluded from
triggering a full restart.
spring.devtools.restart.additional-paths= # Additional paths to watch for changes.
spring.devtools.restart.enabled=true # Enable automatic restart.
spring.devtools.restart.exclude=META-INF/maven/**,META-INF/resources/**,resources/**,static/**,public/
**,templates/**,**/*Test.class,**/*Tests.class,git.properties # Patterns that should be excluded from
triggering a full restart.
spring.devtools.restart.poll-interval=1000 # Amount of time (in milliseconds) to wait between polling
for classpath changes.
spring.devtools.restart.quiet-period=400 # Amount of quiet time (in milliseconds) required without any
classpath changes before a restart is triggered.
spring.devtools.restart.trigger-file= # Name of a specific file that when changed will trigger the
restart check. If not specified any classpath file change will trigger the restart.

# REMOTE DEVTOOLS (RemoteDevToolsProperties)
spring.devtools.remote.context-path=/.~spring-boot!~ # Context path used to handle the remote
connection.
spring.devtools.remote.proxy.host= # The host of the proxy to use to connect to the remote application.
spring.devtools.remote.proxy.port= # The port of the proxy to use to connect to the remote application.
spring.devtools.remote.restart.enabled=true # Enable remote restart.
spring.devtools.remote.secret= # A shared secret required to establish a connection (required to enable
remote support).
spring.devtools.remote.secret-header-name=X-AUTH-TOKEN # HTTP header used to transfer the shared secret.

# -----
# TESTING PROPERTIES
# -----

spring.test.database.replace=any # Type of existing DataSource to replace.
spring.test.mockmvc.print=default # MVC Print option.

```

Appendix B. Configuration meta-data

Spring Boot jars are shipped with meta-data files that provide details of all supported configuration properties. The files are designed to allow IDE developers to offer contextual help and “code completion” as users are working with `application.properties` or `application.yml` files.

The majority of the meta-data file is generated automatically at compile time by processing all items annotated with `@ConfigurationProperties`. However, it is possible to [write part of the meta-data manually](#) for corner cases or more advanced use cases.

B.1 Meta-data format

Configuration meta-data files are located inside jars under `META-INF/spring-configuration-metadata.json`. They use a simple JSON format with items categorized under either “groups” or “properties” and additional values hint categorized under “hints”:

```
{
  "groups": [
    {
      "name": "server",
      "type": "org.springframework.boot.autoconfigure.web.ServerProperties",
      "sourceType": "org.springframework.boot.autoconfigure.web.ServerProperties"
    },
    {
      "name": "spring.jpa.hibernate",
      "type": "org.springframework.boot.autoconfigure.orm.jpa.JpaProperties$Hibernate",
      "sourceType": "org.springframework.boot.autoconfigure.orm.jpa.JpaProperties",
      "sourceMethod": "getHibernate()"
    }
    ...
  ],
  "properties": [
    {
      "name": "server.port",
      "type": "java.lang.Integer",
      "sourceType": "org.springframework.boot.autoconfigure.web.ServerProperties"
    },
    {
      "name": "server.servlet.path",
      "type": "java.lang.String",
      "sourceType": "org.springframework.boot.autoconfigure.web.ServerProperties",
      "defaultValue": "/"
    },
    {
      "name": "spring.jpa.hibernate.ddl-auto",
      "type": "java.lang.String",
      "description": "DDL mode. This is actually a shortcut for the \"hibernate.hbm2ddl.auto\" property.",
      "sourceType": "org.springframework.boot.autoconfigure.orm.jpa.JpaProperties$Hibernate"
    }
    ...
  ],
  "hints": [
    {
      "name": "spring.jpa.hibernate.ddl-auto",
      "values": [
        {
          "value": "none",
          "description": "Disable DDL handling."
        },
        {
          "value": "validate",
          "description": "Validate the schema, make no changes to the database."
        },
        {
          "value": "update",
          "description": "Update the schema if necessary."
        }
      ]
    }
  ]
}
```



```

    },
    {
      "value": "create",
      "description": "Create the schema and destroy previous data."
    },
    {
      "value": "create-drop",
      "description": "Create and then destroy the schema at the end of the session."
    }
  ]
}
}]

```

Each “property” is a configuration item that the user specifies with a given value. For example `server.port` and `server.servlet.path` might be specified in `application.properties` as follows:

```

server.port=9090
server.servlet.path=/home

```

The “groups” are higher level items that don’t themselves specify a value, but instead provide a contextual grouping for properties. For example the `server.port` and `server.servlet.path` properties are part of the `server` group.

Note

It is not required that every “property” has a “group”, some properties might just exist in their own right.

Finally, “hints” are additional information used to assist the user in configuring a given property. When configuring the `spring.jpa.hibernate.ddl-auto` property, a tool can use it to offer some auto-completion help for the `none`, `validate`, `update`, `create` and `create-drop` values.

Group Attributes

The JSON object contained in the `groups` array can contain the following attributes:

Name	Type	Purpose
<code>name</code>	String	The full name of the group. This attribute is mandatory.
<code>type</code>	String	The class name of the data type of the group. For example, if the group was based on a class annotated with <code>@ConfigurationProperties</code> the attribute would contain the fully qualified name of that class. If it was based on a <code>@Bean</code> method, it would be the return type of that method. The attribute may be omitted if the type is not known.
<code>description</code>	String	A short description of the group that can be displayed to users. May be omitted if no description is available. It is recommended that descriptions are a short paragraphs, with the first line providing a concise summary. The last line in the description should end with a period (.).
<code>sourceType</code>	String	The class name of the source that contributed this group. For example, if the group was based on a <code>@Bean</code> method annotated

Name	Type	Purpose
		with <code>@ConfigurationProperties</code> this attribute would contain the fully qualified name of the <code>@Configuration</code> class containing the method. The attribute may be omitted if the source type is not known.
<code>sourceMethod</code>	String	The full name of the method (include parenthesis and argument types) that contributed this group. For example, the name of a <code>@ConfigurationProperties</code> annotated <code>@Bean</code> method. May be omitted if the source method is not known.

Property Attributes

The JSON object contained in the `properties` array can contain the following attributes:

Name	Type	Purpose
<code>name</code>	String	The full name of the property. Names are in lowercase dashed form (e.g. <code>server.servlet.path</code>). This attribute is mandatory.
<code>type</code>	String	The full signature of the data type of the property. For example, <code>java.lang.String</code> but also a full generic type such as <code>java.util.Map<java.util.String,acme.MyEnum></code> . This attribute can be used to guide the user as to the types of values that they can enter. For consistency, the type of a primitive is specified using its wrapper counterpart, i.e. <code>boolean</code> becomes <code>java.lang.Boolean</code> . Note that this class may be a complex type that gets converted from a <code>String</code> as values are bound. May be omitted if the type is not known.
<code>description</code>	String	A short description of the group that can be displayed to users. May be omitted if no description is available. It is recommended that descriptions are a short paragraphs, with the first line providing a concise summary. The last line in the description should end with a period (<code>.</code>).
<code>sourceType</code>	String	The class name of the source that contributed this property. For example, if the property was from a class annotated with <code>@ConfigurationProperties</code> this attribute would contain the fully qualified name of that class. May be omitted if the source type is not known.
<code>defaultValue</code>	Object	The default value which will be used if the property is not specified. Can also be an array of value(s) if the type of the property is an array. May be omitted if the default value is not known.
<code>deprecation</code>	Deprecation	Specify if the property is deprecated. May be omitted if the field is not deprecated or if that information is not known. See below for more details.

The JSON object contained in the `deprecation` attribute of each `properties` element can contain the following attributes:

Name	Type	Purpose
level	String	The level of deprecation, can be either <code>warning</code> (default) or <code>error</code> . When a property has a <code>warning</code> deprecation level it should still be bound in the environment. When it has an <code>error</code> deprecation level however, the property is no longer managed and will not be bound.
reason	String	A short description of the reason why the property was deprecated. May be omitted if no reason is available. It is recommended that descriptions are a short paragraphs, with the first line providing a concise summary. The last line in the description should end with a period (.).
replacement	String	The full name of the property that is <i>replacing</i> this deprecated property. May be omitted if there is no replacement for this property.

Note

Prior to Spring Boot 1.3, a single `deprecated` boolean attribute can be used instead of the deprecation element. This is still supported in a deprecated fashion and should no longer be used. If no reason and replacement are available, an empty `deprecation` object should be set.

Deprecation can also be specified declaratively in code by adding the `@DeprecatedConfigurationProperty` annotation to the getter exposing the deprecated property. For instance, let's assume the `app.foo.target` property was confusing and was renamed to `app.foo.name`

```
@ConfigurationProperties("app.foo")
public class FooProperties {

    private String name;

    public String getName() { ... }

    public void setName(String name) { ... }

    @DeprecatedConfigurationProperty(replacement = "app.foo.name")
    @Deprecated
    public String getTarget() {
        return getName();
    }

    @Deprecated
    public void setTarget(String target) {
        setName(target);
    }
}
```

Note

There is no way to set a `level` as `warning` is always assumed since code is still handling the property.

The code above makes sure that the deprecated property still works (delegating to the `name` property behind the scenes). Once the `getTarget` and `setTarget` methods can be removed from your public

API, the automatic deprecation hint in the meta-data will go away as well. If you want to keep a hint, adding manual meta-data with an `error` deprecation level ensures that users are still informed about that property and is particularly useful when a `replacement` is provided.

Hint Attributes

The JSON object contained in the `hints` array can contain the following attributes:

Name	Type	Purpose
<code>name</code>	String	The full name of the property that this hint refers to. Names are in lowercase dashed form (e.g. <code>server.servlet.path</code>). If the property refers to a map (e.g. <code>system.contexts</code>) the hint either applies to the <i>keys</i> of the map (<code>system.context.keys</code>) or the <i>values</i> (<code>system.context.values</code>). This attribute is mandatory.
<code>values</code>	ValueHint[]	A list of valid values as defined by the <code>ValueHint</code> object (see below). Each entry defines the value and may have a description
<code>providers</code>	ValueProvider[]	A list of providers as defined by the <code>ValueProvider</code> object (see below). Each entry defines the name of the provider and its parameters, if any.

The JSON object contained in the `values` attribute of each `hint` element can contain the following attributes:

Name	Type	Purpose
<code>value</code>	Object	A valid value for the element to which the hint refers to. Can also be an array of value(s) if the type of the property is an array. This attribute is mandatory.
<code>description</code>	String	A short description of the value that can be displayed to users. May be omitted if no description is available. It is recommended that descriptions are a short paragraphs, with the first line providing a concise summary. The last line in the description should end with a period (.

The JSON object contained in the `providers` attribute of each `hint` element can contain the following attributes:

Name	Type	Purpose
<code>name</code>	String	The name of the provider to use to offer additional content assistance for the element to which the hint refers to.
<code>parameters</code>	JSON object	Any additional parameter that the provider supports (check the documentation of the provider for more details).

Repeated meta-data items

It is perfectly acceptable for “property” and “group” objects with the same name to appear multiple times within a meta-data file. For example, you could bind two separate classes to the same prefix, with each

potentially offering overlap of property names. While this is not supposed to be a frequent scenario, consumers of meta-data should take care to ensure that they support such scenarios.

B.2 Providing manual hints

To improve the user experience and further assist the user in configuring a given property, you can provide additional meta-data that:

1. Describes the list of potential values for a property.
2. Associates a provider to attach a well-defined semantic to a property so that a tool can discover the list of potential values based on the project's context.

Value hint

The `name` attribute of each hint refers to the name of a property. In the initial example above, we provide 5 values for the `spring.jpa.hibernate.ddl-auto` property: `none`, `validate`, `update`, `create` and `create-drop`. Each value may have a description as well.

If your property is of type `Map`, you can provide hints for both the keys and the values (but not for the map itself). The special `.keys` and `.values` suffixes must be used to refer to the keys and the values respectively.

Let's assume a `foo.contexts` that maps magic String values to an integer:

```
@ConfigurationProperties("foo")
public class FooProperties {

    private Map<String,Integer> contexts;
    // getters and setters
}
```

The magic values are `foo` and `bar` for instance. In order to offer additional content assistance for the keys, you could add the following to [the manual meta-data of the module](#):

```
{ "hints": [
  {
    "name": "foo.contexts.keys",
    "values": [
      {
        "value": "foo"
      },
      {
        "value": "bar"
      }
    ]
  }
]}
```

Note

Of course, you should have an `Enum` for those two values instead. This is by far the most effective approach to auto-completion if your IDE supports it.

Value provider

Providers are a powerful way of attaching semantics to a property. We define in the section below the official providers that you can use for your own hints. Bare in mind however that your favorite IDE may implement some of these or none of them. It could eventually provide its own as well.

Note

As this is a new feature, IDE vendors will have to catch up with this new feature.

The table below summarizes the list of supported providers:

Name	Description
any	Permit any additional value to be provided.
class-reference	Auto-complete the classes available in the project. Usually constrained by a base class that is specified via the <code>target</code> parameter.
handle-as	Handle the property as if it was defined by the type defined via the mandatory <code>target</code> parameter.
logger-name	Auto-complete valid logger names. Typically, package and class names available in the current project can be auto-completed.
spring-bean-reference	Auto-complete the available bean names in the current project. Usually constrained by a base class that is specified via the <code>target</code> parameter.
spring-profile-name	Auto-complete the available Spring profile names in the project.

Tip

No more than one provider can be active for a given property but you can specify several providers if they can all manage the property *in some ways*. Make sure to place the most powerful provider first as the IDE must use the first one in the JSON section it can handle. If no provider for a given property is supported, no special content assistance is provided either.

Any

The **any** provider permits any additional values to be provided. Regular value validation based on the property type should be applied if this is supported.

This provider will be typically used if you have a list of values and any extra values are still to be considered as valid.

The example below offers `on` and `off` as auto-completion values for `system.state`; any other value is also allowed:

```
{ "hints": [
  {
    "name": "system.state",
    "values": [
      {
        "value": "on"
      },
      {
        "value": "off"
      }
    ],
    "providers": [
      {
        "name": "any"
      }
    ]
  }
]
```

```

    }
  ]
}]

```

Class reference

The **class-reference** provider auto-completes classes available in the project. This provider supports these parameters:

Parameter	Type	Default value	Description
target	String (Class)	<i>none</i>	The fully qualified name of the class that should be assignable to the chosen value. Typically used to filter out non candidate classes. Note that this information can be provided by the type itself by exposing a class with the appropriate upper bound.
concrete	boolean	true	Specify if only concrete classes are to be considered as valid candidates.

The meta-data snippet below corresponds to the standard `server.servlet.jsp.class-name` property that defines the `JspServlet` class name to use:

```

{
  "hints": [
    {
      "name": "server.servlet.jsp.class-name",
      "providers": [
        {
          "name": "class-reference",
          "parameters": {
            "target": "javax.servlet.http.HttpServlet"
          }
        }
      ]
    }
  ]
}

```

Handle As

The **handle-as** provider allows you to substitute the type of the property to a more high-level type. This typically happens when the property has a `java.lang.String` type because you don't want your configuration classes to rely on classes that may not be on the classpath. This provider supports these parameters:

Parameter	Type	Default value	Description
target	String (Class)	<i>none</i>	The fully qualified name of the type to consider for the property. This parameter is mandatory.

The following types can be used:

- Any `java.lang.Enum` that lists the possible values for the property (By all means, try to define the property with the `Enum` type instead as no further hint should be required for the IDE to auto-complete the values).

- `java.nio.charset.Charset`: auto-completion of charset/encoding values (e.g. UTF-8)
- `java.util.Locale`: auto-completion of locales (e.g. en_US)
- `org.springframework.util.MimeType`: auto-completion of content type values (e.g. text/plain)
- `org.springframework.core.io.Resource`: auto-completion of Spring's Resource abstraction to refer to a file on the filesystem or on the classpath. (e.g. `classpath:/foo.properties`)

Note

If multiple values can be provided, use a `Collection` or `Array` type to teach the IDE about it.

The meta-data snippet below corresponds to the standard `liquibase.change-log` property that defines the path to the changelog to use. It is actually used internally as a `org.springframework.core.io.Resource` but cannot be exposed as such as we need to keep the original String value to pass it to the Liquibase API.

```
{
  "hints": [
    {
      "name": "liquibase.change-log",
      "providers": [
        {
          "name": "handle-as",
          "parameters": {
            "target": "org.springframework.core.io.Resource"
          }
        }
      ]
    }
  ]
}
```

Logger name

The **logger-name** provider auto-completes valid logger names. Typically, package and class names available in the current project can be auto-completed. Specific frameworks may have extra magic logger names that could be supported as well.

Since a logger name can be any arbitrary name, really, this provider should allow any value but could highlight valid packages and class names that are not available in the project's classpath.

The meta-data snippet below corresponds to the standard `logging.level` property, keys are *logger names* and values correspond to the standard log levels or any custom level:

```
{
  "hints": [
    {
      "name": "logging.level.keys",
      "values": [
        {
          "value": "root",
          "description": "Root logger used to assign the default logging level."
        }
      ],
      "providers": [
        {
          "name": "logger-name"
        }
      ]
    }
  ],
  {
```



```

    "name": "logging.level.values",
    "values": [
      {
        "value": "trace"
      },
      {
        "value": "debug"
      },
      {
        "value": "info"
      },
      {
        "value": "warn"
      },
      {
        "value": "error"
      },
      {
        "value": "fatal"
      },
      {
        "value": "off"
      }
    ],
    "providers": [
      {
        "name": "any"
      }
    ]
  }
}

```

Spring bean reference

The **spring-bean-reference** provider auto-completes the beans that are defined in the configuration of the current project. This provider supports these parameters:

Parameter	Type	Default value	Description
target	String (Class)	<i>none</i>	The fully qualified name of the bean class that should be assignable to the candidate. Typically used to filter out non candidate beans.

The meta-data snippet below corresponds to the standard `spring.jmx.server` property that defines the name of the `MBeanServer` bean to use:

```

{
  "hints": [
    {
      "name": "spring.jmx.server",
      "providers": [
        {
          "name": "spring-bean-reference",
          "parameters": {
            "target": "javax.management.MBeanServer"
          }
        }
      ]
    }
  ]
}

```

Note

The binder is not aware of the meta-data so if you provide that hint, you will still need to transform the bean name into an actual Bean reference using the `ApplicationContext`.

Spring profile name

The `spring-profile-name` provider auto-completes the Spring profiles that are defined in the configuration of the current project.

The meta-data snippet below corresponds to the standard `spring.profiles.active` property that defines the name of the Spring profile(s) to enable:

```
{ "hints": [
  {
    "name": "spring.profiles.active",
    "providers": [
      {
        "name": "spring-profile-name"
      }
    ]
  }
]}
```

B.3 Generating your own meta-data using the annotation processor

You can easily generate your own configuration meta-data file from items annotated with `@ConfigurationProperties` by using the `spring-boot-configuration-processor` jar. The jar includes a Java annotation processor which is invoked as your project is compiled. To use the processor, simply include `spring-boot-configuration-processor` as an optional dependency, for example with Maven you would add:

```
<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-configuration-processor</artifactId>
  <optional>true</optional>
</dependency>
```

With Gradle, you can use the [propdeps-plugin](#) and specify:

```
dependencies {
  optional "org.springframework.boot:spring-boot-configuration-processor"
}

compileJava.dependsOn(processResources)
```

Note

You need to add `compileJava.dependsOn(processResources)` to your build to ensure that resources are processed before code is compiled. Without this directive any additional `spring-configuration-metadata.json` files will not be processed.

The processor will pick up both classes and methods that are annotated with `@ConfigurationProperties`. The Javadoc for field values within configuration classes will be used to populate the `description` attribute.

Note

You should only use simple text with `@ConfigurationProperties` field Javadoc since they are not processed before being added to the JSON.

Properties are discovered via the presence of standard getters and setters with special handling for collection types (that will be detected even if only a getter is present). The annotation processor also supports the use of the `@Data`, `@Getter` and `@Setter` lombok annotations.

Note

If you are using AspectJ in your project, you need to make sure that the annotation processor only runs once. There are several ways to do this: with Maven, you can configure the `maven-apt-plugin` explicitly and add the dependency to the annotation processor only there. You could also let the AspectJ plugin run all the processing and disable annotation processing in the `maven-compiler-plugin` configuration:

```
<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-compiler-plugin</artifactId>
  <configuration>
    <proc>none</proc>
  </configuration>
</plugin>
```

Nested properties

The annotation processor will automatically consider inner classes as nested properties. For example, the following class:

```
@ConfigurationProperties(prefix="server")
public class ServerProperties {

    private String name;

    private Host host;

    // ... getter and setters

    private static class Host {

        private String ip;

        private int port;

        // ... getter and setters

    }

}
```

Will produce meta-data information for `server.name`, `server.host.ip` and `server.host.port` properties. You can use the `@NestedConfigurationProperty` annotation on a field to indicate that a regular (non-inner) class should be treated as if it were nested.

Adding additional meta-data

Spring Boot's configuration file handling is quite flexible; and it is often the case that properties may exist that are not bound to a `@ConfigurationProperties` bean. You may also need to tune some attributes of an existing key. To support such cases and allow you to provide custom "hints", the annotation processor will automatically merge items from `META-INF/additional-spring-configuration-metadata.json` into the main meta-data file.

If you refer to a property that has been detected automatically, the description, default value and deprecation information are overridden if specified. If the manual property declaration is not identified in the current module, it is added as a brand new property.

The format of the `additional-spring-configuration-metadata.json` file is exactly the same as the regular `spring-configuration-metadata.json`. The additional properties file is optional, if you don't have any additional properties, simply don't add it.

Appendix C. Auto-configuration classes

Here is a list of all auto-configuration classes provided by Spring Boot with links to documentation and source code. Remember to also look at the autoconfig report in your application for more details of which features are switched on. (start the app with `--debug` or `-Ddebug`, or in an Actuator application use the `autoconfig` endpoint).

C.1 From the “spring-boot-autoconfigure” module

The following auto-configuration classes are from the `spring-boot-autoconfigure` module:

Configuration Class	Links
ActiveMQAutoConfiguration	javadoc
AopAutoConfiguration	javadoc
ArtemisAutoConfiguration	javadoc
BatchAutoConfiguration	javadoc
CacheAutoConfiguration	javadoc
CassandraAutoConfiguration	javadoc
CassandraDataAutoConfiguration	javadoc
CassandraReactiveDataAutoConfiguration	javadoc
CassandraReactiveRepositoriesAutoConfiguration	javadoc
CassandraRepositoriesAutoConfiguration	javadoc
CloudAutoConfiguration	javadoc
CodecsAutoConfiguration	javadoc
ConfigurationPropertiesAutoConfiguration	javadoc
CouchbaseAutoConfiguration	javadoc
CouchbaseDataAutoConfiguration	javadoc
CouchbaseRepositoriesAutoConfiguration	javadoc
DataSourceAutoConfiguration	javadoc
DataSourceTransactionManagerAutoConfiguration	javadoc
DeviceDelegatingViewResolverAutoConfiguration	javadoc
DeviceResolverAutoConfiguration	javadoc
DispatcherServletAutoConfiguration	javadoc
ElasticsearchAutoConfiguration	javadoc

Configuration Class	Links
ElasticsearchDataAutoConfiguration	javadoc
ElasticsearchRepositoriesAutoConfiguration	javadoc
EmbeddedLdapAutoConfiguration	javadoc
EmbeddedMongoAutoConfiguration	javadoc
ErrorMvcAutoConfiguration	javadoc
FacebookAutoConfiguration	javadoc
FallbackWebSecurityAutoConfiguration	javadoc
FlywayAutoConfiguration	javadoc
FreeMarkerAutoConfiguration	javadoc
GroovyTemplateAutoConfiguration	javadoc
GsonAutoConfiguration	javadoc
H2ConsoleAutoConfiguration	javadoc
HazelcastAutoConfiguration	javadoc
HazelcastJpaDependencyAutoConfiguration	javadoc
HibernateJpaAutoConfiguration	javadoc
HttpEncodingAutoConfiguration	javadoc
HttpHandlerAutoConfiguration	javadoc
HttpMessageConvertersAutoConfiguration	javadoc
HypermediaAutoConfiguration	javadoc
InfluxDbAutoConfiguration	javadoc
IntegrationAutoConfiguration	javadoc
JacksonAutoConfiguration	javadoc
JdbcTemplateAutoConfiguration	javadoc
JerseyAutoConfiguration	javadoc
JestAutoConfiguration	javadoc
JmsAutoConfiguration	javadoc
JmxAutoConfiguration	javadoc
JndiConnectionFactoryAutoConfiguration	javadoc
JndiDataSourceAutoConfiguration	javadoc
JooqAutoConfiguration	javadoc

Configuration Class	Links
JpaRepositoriesAutoConfiguration	javadoc
JtaAutoConfiguration	javadoc
KafkaAutoConfiguration	javadoc
LdapAutoConfiguration	javadoc
LdapDataAutoConfiguration	javadoc
LdapRepositoriesAutoConfiguration	javadoc
LinkedInAutoConfiguration	javadoc
LiquibaseAutoConfiguration	javadoc
MailSenderAutoConfiguration	javadoc
MailSenderValidatorAutoConfiguration	javadoc
MessageSourceAutoConfiguration	javadoc
MongoAutoConfiguration	javadoc
MongoDataAutoConfiguration	javadoc
MongoReactiveAutoConfiguration	javadoc
MongoReactiveDataAutoConfiguration	javadoc
MongoReactiveRepositoriesAutoConfiguration	javadoc
MongoRepositoriesAutoConfiguration	javadoc
MultipartAutoConfiguration	javadoc
MustacheAutoConfiguration	javadoc
Neo4jDataAutoConfiguration	javadoc
Neo4jRepositoriesAutoConfiguration	javadoc
OAuth2AutoConfiguration	javadoc
PersistenceExceptionTranslationAutoConfiguration	javadoc
ProjectInfoAutoConfiguration	javadoc
PropertyPlaceholderAutoConfiguration	javadoc
QuartzAutoConfiguration	javadoc
RabbitAutoConfiguration	javadoc
ReactiveWebServerAutoConfiguration	javadoc
ReactorCoreAutoConfiguration	javadoc
RedisAutoConfiguration	javadoc

Configuration Class	Links
RedisReactiveAutoConfiguration	javadoc
RedisRepositoriesAutoConfiguration	javadoc
RepositoryRestMvcAutoConfiguration	javadoc
RestTemplateAutoConfiguration	javadoc
SecurityAutoConfiguration	javadoc
SecurityFilterAutoConfiguration	javadoc
SendGridAutoConfiguration	javadoc
ServletWebServerFactoryAutoConfiguration	javadoc
SessionAutoConfiguration	javadoc
SitePreferenceAutoConfiguration	javadoc
SocialWebAutoConfiguration	javadoc
SolrAutoConfiguration	javadoc
SolrRepositoriesAutoConfiguration	javadoc
SpringApplicationAdminJmxAutoConfiguration	javadoc
SpringDataWebAutoConfiguration	javadoc
ThymeleafAutoConfiguration	javadoc
TransactionAutoConfiguration	javadoc
TwitterAutoConfiguration	javadoc
ValidationAutoConfiguration	javadoc
WebClientAutoConfiguration	javadoc
WebFluxAutoConfiguration	javadoc
WebMvcAutoConfiguration	javadoc
WebServicesAutoConfiguration	javadoc
WebSocketMessagingAutoConfiguration	javadoc
WebSocketReactiveAutoConfiguration	javadoc
WebSocketServletAutoConfiguration	javadoc
XADataSourceAutoConfiguration	javadoc

C.2 From the “spring-boot-actuator” module

The following auto-configuration classes are from the `spring-boot-actuator` module:

Configuration Class	Links
AuditAutoConfiguration	javadoc
CacheStatisticsAutoConfiguration	javadoc
CloudFoundryActuatorAutoConfiguration	javadoc
EndpointAutoConfiguration	javadoc
EndpointMBeanExportAutoConfiguration	javadoc
EndpointWebMvcAutoConfiguration	javadoc
HealthIndicatorAutoConfiguration	javadoc
InfoContributorAutoConfiguration	javadoc
JolokiaAutoConfiguration	javadoc
ManagementWebSecurityAutoConfiguration	javadoc
MetricExportAutoConfiguration	javadoc
MetricFilterAutoConfiguration	javadoc
MetricRepositoryAutoConfiguration	javadoc
MetricsChannelAutoConfiguration	javadoc
MetricsDropwizardAutoConfiguration	javadoc
PublicMetricsAutoConfiguration	javadoc
TraceRepositoryAutoConfiguration	javadoc
TraceWebFilterAutoConfiguration	javadoc

Appendix D. Test auto-configuration annotations

Here is a table of the various `@...Test` annotations that can be used to test slices of your application and the auto-configuration that they import by default:

Test slice	Imported auto-configuration
<code>@DataJpaTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.data.jpa.J org.springframework.boot.autoconfigure.flyway.Fly org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jdbc.JdbcT org.springframework.boot.autoconfigure.liquibase. org.springframework.boot.autoconfigure.orm.jpa.Hi org.springframework.boot.autoconfigure.transactio org.springframework.boot.test.autoconfigure.jdbc. org.springframework.boot.test.autoconfigure.orm.j</pre>
<code>@DataLdapTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.data.ldap. org.springframework.boot.autoconfigure.data.ldap. org.springframework.boot.autoconfigure.ldap.LdapA org.springframework.boot.autoconfigure.ldap.embed</pre>
<code>@DataMongoTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.data.mongo org.springframework.boot.autoconfigure.data.mongo org.springframework.boot.autoconfigure.data.mongo org.springframework.boot.autoconfigure.mongo.Mong org.springframework.boot.autoconfigure.mongo.Mong org.springframework.boot.autoconfigure.mongo.embe</pre>
<code>@DataNeo4jTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.data.neo4j org.springframework.boot.autoconfigure.data.neo4j org.springframework.boot.autoconfigure.transactio</pre>
<code>@DataRedisTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.data.redis org.springframework.boot.autoconfigure.data.redis</pre>
<code>@JdbcTest</code>	<pre>org.springframework.boot.autoconfigure.cache.Cache org.springframework.boot.autoconfigure.flyway.Fly org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jdbc.JdbcT org.springframework.boot.autoconfigure.liquibase.</pre>

Test slice	Imported auto-configuration
	<pre>org.springframework.boot.autoconfigure.transaction org.springframework.boot.test.autoconfigure.jdbc.</pre>
@JooqTest	<pre>org.springframework.boot.autoconfigure.flyway.Fly org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jdbc.DataS org.springframework.boot.autoconfigure.jooq.JooqA org.springframework.boot.autoconfigure.liquibase. org.springframework.boot.autoconfigure.transaction</pre>
@JsonTest	<pre>org.springframework.boot.autoconfigure.cache.Cach org.springframework.boot.autoconfigure.gson.GsonA org.springframework.boot.autoconfigure.jackson.Ja org.springframework.boot.test.autoconfigure.json.</pre>
@RestClientTest	<pre>org.springframework.boot.autoconfigure.cache.Cach org.springframework.boot.autoconfigure.gson.GsonA org.springframework.boot.autoconfigure.http.HttpM org.springframework.boot.autoconfigure.http.codec org.springframework.boot.autoconfigure.jackson.Ja org.springframework.boot.autoconfigure.web.client org.springframework.boot.autoconfigure.web.reacti org.springframework.boot.test.autoconfigure.web.c org.springframework.boot.test.autoconfigure.web.c</pre>
@WebFluxTest	<pre>org.springframework.boot.autoconfigure.cache.Cach org.springframework.boot.autoconfigure.context.Me org.springframework.boot.autoconfigure.validation org.springframework.boot.autoconfigure.web.reacti org.springframework.boot.test.autoconfigure.web.r</pre>
@WebMvcTest	<pre>org.springframework.boot.autoconfigure.cache.Cach org.springframework.boot.autoconfigure.context.Me org.springframework.boot.autoconfigure.freemarker org.springframework.boot.autoconfigure.groovy.tem org.springframework.boot.autoconfigure.gson.GsonA org.springframework.boot.autoconfigure.hateoas.Hy org.springframework.boot.autoconfigure.http.HttpM org.springframework.boot.autoconfigure.jackson.Ja org.springframework.boot.autoconfigure.mustache.M org.springframework.boot.autoconfigure.thymeleaf. org.springframework.boot.autoconfigure.validation org.springframework.boot.autoconfigure.web.servle org.springframework.boot.autoconfigure.web.servle org.springframework.boot.test.autoconfigure.web.s org.springframework.boot.test.autoconfigure.web.s org.springframework.boot.test.autoconfigure.web.s org.springframework.boot.test.autoconfigure.web.s</pre>

Appendix E. The executable jar format

The `spring-boot-loader` modules allows Spring Boot to support executable jar and war files. If you're using the Maven or Gradle plugin, executable jars are automatically generated and you generally won't need to know the details of how they work.

If you need to create executable jars from a different build system, or if you are just curious about the underlying technology, this section provides some background.

E.1 Nested JARs

Java does not provide any standard way to load nested jar files (i.e. jar files that are themselves contained within a jar). This can be problematic if you are looking to distribute a self-contained application that you can just run from the command line without unpacking.

To solve this problem, many developers use "shaded" jars. A shaded jar simply packages all classes, from all jars, into a single 'uber jar'. The problem with shaded jars is that it becomes hard to see which libraries you are actually using in your application. It can also be problematic if the same filename is used (but with different content) in multiple jars. Spring Boot takes a different approach and allows you to actually nest jars directly.

The executable jar file structure

Spring Boot Loader compatible jar files should be structured in the following way:

```
example.jar
|
+-META-INF
| +-MANIFEST.MF
+-org
| +-springframework
| | +-boot
| | | +-loader
| | | +-<spring boot loader classes>
+-BOOT-INF
| +-classes
| | +-mycompany
| | | +-project
| | | +-YourClasses.class
+-lib
| +-dependency1.jar
| +-dependency2.jar
```

Application classes should be placed in a nested `BOOT-INF/classes` directory. Dependencies should be placed in a nested `BOOT-INF/lib` directory.

The executable war file structure

Spring Boot Loader compatible war files should be structured in the following way:

```
example.war
|
+-META-INF
| +-MANIFEST.MF
+-org
```

```

| +-springframework
| | +-boot
| | | +-loader
| | | +-<spring boot loader classes>
+-WEB-INF
  +-classes
  | +-com
  | | +-mycompany
  | | | +-project
  | | | | +-YourClasses.class
  +-lib
  | +-dependency1.jar
  | +-dependency2.jar
+-lib-provided
  +-servlet-api.jar
  +-dependency3.jar

```

Dependencies should be placed in a nested `WEB-INF/lib` directory. Any dependencies that are required when running embedded but are not required when deploying to a traditional web container should be placed in `WEB-INF/lib-provided`.

E.2 Spring Boot’s “JarFile” class

The core class used to support loading nested jars is `org.springframework.boot.loader.jar.JarFile`. It allows you to load jar content from a standard jar file, or from nested child jar data. When first loaded, the location of each `JarEntry` is mapped to a physical file offset of the outer jar:

```

myapp.jar
+-----+
| /BOOT-INF/classes | /BOOT-INF/lib/mylib.jar | | | | | |
|+-----+|+-----+|
||   A.class   ||| B.class | C.class ||
|+-----+|+-----+|
+-----+
^           ^           ^
0063           3452       3980

```

The example above shows how `A.class` can be found in `/BOOT-INF/classes` in `myapp.jar` position 0063. `B.class` from the nested jar can actually be found in `myapp.jar` position 3452 and `C.class` is at position 3980.

Armed with this information, we can load specific nested entries by simply seeking to the appropriate part of the outer jar. We don’t need to unpack the archive and we don’t need to read all entry data into memory.

Compatibility with the standard Java “JarFile”

Spring Boot Loader strives to remain compatible with existing code and libraries. `org.springframework.boot.loader.jar.JarFile` extends from `java.util.jar.JarFile` and should work as a drop-in replacement. The `getURL()` method will return a URL that opens a `java.net.JarURLConnection` compatible connection and can be used with Java’s `URLClassLoader`.

E.3 Launching executable jars

The `org.springframework.boot.loader.Launcher` class is a special bootstrap class that is used as an executable jars main entry point. It is the actual `Main-Class` in your jar file and it’s used to setup an appropriate `URLClassLoader` and ultimately call your `main()` method.

There are 3 launcher subclasses (`JarLauncher`, `WarLauncher` and `PropertiesLauncher`). Their purpose is to load resources (`.class` files etc.) from nested jar files or war files in directories (as opposed to explicitly on the classpath). In the case of `JarLauncher` and `WarLauncher` the nested paths are fixed. `JarLauncher` looks in `BOOT-INF/lib/` and `WarLauncher` looks in `WEB-INF/lib/` and `WEB-INF/lib-provided/` so you just add extra jars in those locations if you want more. The `PropertiesLauncher` looks in `BOOT-INF/lib/` in your application archive by default, but you can add additional locations by setting an environment variable `LOADER_PATH` or `loader.path` in `loader.properties` (comma-separated list of directories, archives, or directories within archives).

Launcher manifest

You need to specify an appropriate Launcher as the `Main-Class` attribute of `META-INF/MANIFEST.MF`. The actual class that you want to launch (i.e. the class that you wrote that contains a main method) should be specified in the `Start-Class` attribute.

For example, here is a typical `MANIFEST.MF` for an executable jar file:

```
Main-Class: org.springframework.boot.loader.JarLauncher
Start-Class: com.mycompany.project.MyApplication
```

For a war file, it would be:

```
Main-Class: org.springframework.boot.loader.WarLauncher
Start-Class: com.mycompany.project.MyApplication
```

Note

You do not need to specify `Class-Path` entries in your manifest file, the classpath will be deduced from the nested jars.

Exploded archives

Certain PaaS implementations may choose to unpack archives before they run. For example, Cloud Foundry operates in this way. You can run an unpacked archive by simply starting the appropriate launcher:

```
$ unzip -q myapp.jar
$ java org.springframework.boot.loader.JarLauncher
```

E.4 PropertiesLauncher Features

`PropertiesLauncher` has a few special features that can be enabled with external properties (System properties, environment variables, manifest entries or `loader.properties`).

Key	Purpose
<code>loader.path</code>	Comma-separated Classpath, e.g. <code>lib</code> , <code>\${HOME}/app/lib</code> . Earlier entries take precedence, just like a regular <code>-classpath</code> on the <code>javac</code> command line.
<code>loader.home</code>	Used to resolve relative paths in <code>loader.path</code> . E.g. <code>loader.path=lib</code> then <code>\${loader.home}/lib</code> is a classpath location

Key	Purpose
	(along with all jar files in that directory). Also used to locate a <code>loader.properties</code> file. Example <code>/opt/app</code> (defaults to <code>\${user.dir}</code>).
<code>loader.args</code>	Default arguments for the main method (space separated)
<code>loader.main</code>	Name of main class to launch, e.g. <code>com.app.Application</code> .
<code>loader.config.name</code>	Name of properties file, e.g. <code>launcher</code> (defaults to <code>loader</code>).
<code>loader.config.location</code>	Path to properties file, e.g. <code>classpath:loader.properties</code> (defaults to <code>loader.properties</code>).
<code>loader.system</code>	Boolean flag to indicate that all properties should be added to System properties (defaults to <code>false</code>)

When specified as environment variables or manifest entries, the following names should be used:

Key	Manifest entry	Environment variable
<code>loader.path</code>	<code>Loader-Path</code>	<code>LOADER_PATH</code>
<code>loader.home</code>	<code>Loader-Home</code>	<code>LOADER_HOME</code>
<code>loader.args</code>	<code>Loader-Args</code>	<code>LOADER_ARGS</code>
<code>loader.main</code>	<code>Start-Class</code>	<code>LOADER_MAIN</code>
<code>loader.config.location</code>	<code>Loader-Config-Location</code>	<code>LOADER_CONFIG_LOCATION</code>
<code>loader.system</code>	<code>Loader-System</code>	<code>LOADER_SYSTEM</code>

Tip

Build plugins automatically move the `Main-Class` attribute to `Start-Class` when the fat jar is built. If you are using that, specify the name of the class to launch using the `Main-Class` attribute and leave out `Start-Class`.

- `loader.properties` are searched for in `loader.home` then in the root of the classpath, then in `classpath:/BOOT-INF/classes`. The first location that exists is used.
- `loader.home` is only the directory location of an additional properties file (overriding the default) as long as `loader.config.location` is not specified.
- `loader.path` can contain directories (scanned recursively for jar and zip files), archive paths, a directory within an archive that is scanned for jar files (for example, `dependencies.jar!/lib`), or wildcard patterns (for the default JVM behavior). Archive paths can be relative to `loader.home`, or anywhere in the file system with a `jar:file:` prefix.

- `loader.path` (if empty) defaults to `BOOT-INF/lib` (meaning a local directory or a nested one if running from an archive). Because of this `PropertiesLauncher` behaves the same as `JarLauncher` when no additional configuration is provided.
- `loader.path` can not be used to configure the location of `loader.properties` (the classpath used to search for the latter is the JVM classpath when `PropertiesLauncher` is launched).
- Placeholder replacement is done from System and environment variables plus the properties file itself on all values before use.
- The search order for properties (where it makes sense to look in more than one place) is env vars, system properties, `loader.properties`, exploded archive manifest, archive manifest.

E.5 Executable jar restrictions

There are a number of restrictions that you need to consider when working with a Spring Boot Loader packaged application.

Zip entry compression

The `ZipEntry` for a nested jar must be saved using the `ZipEntry.STORED` method. This is required so that we can seek directly to individual content within the nested jar. The content of the nested jar file itself can still be compressed, as can any other entries in the outer jar.

System ClassLoader

Launched applications should use `Thread.getContextClassLoader()` when loading classes (most libraries and frameworks will do this by default). Trying to load nested jar classes via `ClassLoader.getSystemClassLoader()` will fail. Please be aware that `java.util.Logging` always uses the system classloader, for this reason you should consider a different logging implementation.

E.6 Alternative single jar solutions

If the above restrictions mean that you cannot use Spring Boot Loader the following alternatives could be considered:

- [Maven Shade Plugin](#)
- [JarClassLoader](#)
- [OneJar](#)

Appendix F. Dependency versions

The table below provides details of all of the dependency versions that are provided by Spring Boot in its CLI, Maven dependency management and Gradle plugin. When you declare a dependency on one of these artifacts without declaring a version the version that is listed in the table will be used.

Group ID	Artifact ID	Version
antlr	antlr	2.7.7
ch.qos.logback	logback-access	1.2.3
ch.qos.logback	logback-classic	1.2.3
ch.qos.logback	logback-core	1.2.3
com.atomikos	transactions-jdbc	4.0.4
com.atomikos	transactions-jms	4.0.4
com.atomikos	transactions-jta	4.0.4
com.couchbase.client	couchbase-spring-cache	2.1.0
com.couchbase.client	java-client	2.4.7
com.datastax.cassandra	cassandra-driver-core	3.3.0
com.datastax.cassandra	cassandra-driver-mapping	3.3.0
com.fasterxml	classmate	1.3.3
com.fasterxml.jackson.core	jackson-annotations	2.9.0.pr4
com.fasterxml.jackson.core	jackson-core	2.9.0.pr4
com.fasterxml.jackson.core	jackson-databind	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-avro	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-cbor	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-csv	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-ion	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-properties	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-protobuf	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-smile	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-xml	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-dataformat-yaml	2.9.0.pr4

Group ID	Artifact ID	Version
com.fasterxml.jackson.datatype	jackson-datatype-guava	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-hibernate3	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-hibernate4	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-hibernate5	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-hppc	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-jaxrs	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-jdk8	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-joda	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-json-org	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-jsr310	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-jsr353	2.9.0.pr4
com.fasterxml.jackson.datatype	jackson-datatype-pcollections	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-base	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-cbor-provider	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-json-provider	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-smile-provider	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-xml-provider	2.9.0.pr4
com.fasterxml.jackson.jaxrs	jackson-jaxrs-yaml-provider	2.9.0.pr4
com.fasterxml.jackson.jr	jackson-jr-all	2.9.0.pr4
com.fasterxml.jackson.jr	jackson-jr-objects	2.9.0.pr4
com.fasterxml.jackson.jr	jackson-jr-retrofit2	2.9.0.pr4
com.fasterxml.jackson.jr	jackson-jr-stree	2.9.0.pr4
com.fasterxml.jackson.module	jackson-module-afterburner	2.9.0.pr4

Group ID	Artifact ID	Version
com.fasterxml.jackson.module-guice	jackson-module-guice	2.9.0.pr4
com.fasterxml.jackson.module-jaxb-annotations	jackson-module-jaxb-annotations	2.9.0.pr4
com.fasterxml.jackson.module-jsonSchema	jackson-module-jsonSchema	2.9.0.pr4
com.fasterxml.jackson.module-kotlin	jackson-module-kotlin	2.9.0.pr4
com.fasterxml.jackson.module-mrbean	jackson-module-mrbean	2.9.0.pr4
com.fasterxml.jackson.module-osgi	jackson-module-osgi	2.9.0.pr4
com.fasterxml.jackson.module-parameter-names	jackson-module-parameter-names	2.9.0.pr4
com.fasterxml.jackson.module-paranamer	jackson-module-paranamer	2.9.0.pr4
com.fasterxml.jackson.module-scala_2.10	jackson-module-scala_2.10	2.9.0.pr4
com.fasterxml.jackson.module-scala_2.11	jackson-module-scala_2.11	2.9.0.pr4
com.fasterxml.jackson.module-scala_2.12	jackson-module-scala_2.12	2.9.0.pr4
com.github.ben-manes.caffeine	caffeine	2.5.3
com.github.mxab.thymeleaf-extras-data-attribute	thymeleaf-extras-data-attribute	2.0.1
com.google.appengine	appengine-api-1.0-sdk	1.9.54
com.google.code.gson	gson	2.8.1
com.googlecode.json-simple	json-simple	1.1.1
com.h2database	h2	1.4.196
com.hazelcast	hazelcast	3.8.3
com.hazelcast	hazelcast-client	3.8.3
com.hazelcast	hazelcast-hibernate52	1.2.1
com.hazelcast	hazelcast-spring	3.8.3
com.jayway.jsonpath	json-path	2.4.0
com.jayway.jsonpath	json-path-assert	2.4.0
com.microsoft.sqlserver	mssql-jdbc	6.2.1.jre8

Group ID	Artifact ID	Version
com.querydsl	querydsl-apt	4.1.4
com.querydsl	querydsl-collections	4.1.4
com.querydsl	querydsl-core	4.1.4
com.querydsl	querydsl-jpa	4.1.4
com.querydsl	querydsl-mongodb	4.1.4
com.samskivert	jmustache	1.13
com.sendgrid	sendgrid-java	4.0.1
com.sun.mail	javax.mail	1.6.0
com.timgroup	java-statsd-client	3.1.0
com.unboundid	unboundid-ldapsdk	3.2.1
com.zaxxer	HikariCP	2.6.3
commons-beanutils	commons-beanutils	1.9.3
commons-codec	commons-codec	1.10
commons-collections	commons-collections	3.2.2
commons-digester	commons-digester	2.1
commons-pool	commons-pool	1.6
de.flapdoodle.embed	de.flapdoodle.embed.mongo	2.0.0
dom4j	dom4j	1.6.1
io.dropwizard.metrics	metrics-core	3.2.3
io.dropwizard.metrics	metrics-ganglia	3.2.3
io.dropwizard.metrics	metrics-graphite	3.2.3
io.dropwizard.metrics	metrics-servlets	3.2.3
io.lettuce	lettuce-core	5.0.0.RC1
io.netty	netty-all	4.1.13.Final
io.netty	netty-buffer	4.1.13.Final
io.netty	netty-codec	4.1.13.Final
io.netty	netty-codec-dns	4.1.13.Final
io.netty	netty-codec-haproxy	4.1.13.Final
io.netty	netty-codec-http	4.1.13.Final
io.netty	netty-codec-http2	4.1.13.Final

Group ID	Artifact ID	Version
io.netty	netty-codec-memcache	4.1.13.Final
io.netty	netty-codec-mqtt	4.1.13.Final
io.netty	netty-codec-redis	4.1.13.Final
io.netty	netty-codec-smtp	4.1.13.Final
io.netty	netty-codec-socks	4.1.13.Final
io.netty	netty-codec-stomp	4.1.13.Final
io.netty	netty-codec-xml	4.1.13.Final
io.netty	netty-common	4.1.13.Final
io.netty	netty-dev-tools	4.1.13.Final
io.netty	netty-example	4.1.13.Final
io.netty	netty-handler	4.1.13.Final
io.netty	netty-handler-proxy	4.1.13.Final
io.netty	netty-resolver	4.1.13.Final
io.netty	netty-resolver-dns	4.1.13.Final
io.netty	netty-transport	4.1.13.Final
io.netty	netty-transport-native-epoll	4.1.13.Final
io.netty	netty-transport-native-kqueue	4.1.13.Final
io.netty	netty-transport-native-unix-common	4.1.13.Final
io.netty	netty-transport-rxtx	4.1.13.Final
io.netty	netty-transport-sctp	4.1.13.Final
io.netty	netty-transport-udt	4.1.13.Final
io.netty	netty-transport-unix-common	4.1.13.Final-SNAPSHOT
io.projectreactor	reactor-core	3.1.0.M3
io.projectreactor	reactor-test	3.1.0.M3
io.projectreactor.addons	reactor-adapter	3.1.0.M3
io.projectreactor.addons	reactor-extra	3.1.0.M3
io.projectreactor.addons	reactor-logback	3.1.0.M3
io.projectreactor.ipc	reactor-ipc	0.7.0.M1

Group ID	Artifact ID	Version
io.projectreactor.ipc	reactor-netty	0.7.0.M1
io.projectreactor.kafka	reactor-kafka	1.0.0.M3
io.reactivex	rxjava	1.3.0
io.reactivex	rxjava-reactive-streams	1.2.1
io.reactivex.rxjava2	rxjava	2.1.1
io.searchbox	jest	5.3.2
io.undertow	undertow-core	1.4.18.Final
io.undertow	undertow-servlet	1.4.18.Final
io.undertow	undertow-websockets-jsr	1.4.18.Final
javax.cache	cache-api	1.0.0
javax.jms	javax.jms-api	2.0.1
javax.mail	javax.mail-api	1.6.0
javax.servlet	javax.servlet-api	3.1.0
javax.servlet	jstl	1.2
javax.transaction	javax.transaction-api	1.2
javax.validation	validation-api	1.1.0.Final
jaxen	jaxen	1.1.6
joda-time	joda-time	2.9.9
junit	junit	4.12
mysql	mysql-connector-java	5.1.43
net.bytebuddy	byte-buddy	1.6.14
net.bytebuddy	byte-buddy-agent	1.6.14
net.java.dev.jna	jna	4.4.0
net.sf.ehcache	ehcache	2.10.4
net.sourceforge.htmlunit	htmlunit	2.27
net.sourceforge.jtlds	jtlds	1.3.1
net.sourceforge.nekohtml	nekohtml	1.9.22
nz.net.ultraq.thymeleaf	thymeleaf-layout-dialect	2.2.2
org.apache.activemq	activemq-amqp	5.15.0
org.apache.activemq	activemq-blueprint	5.15.0

Group ID	Artifact ID	Version
org.apache.activemq	activemq-broker	5.15.0
org.apache.activemq	activemq-camel	5.15.0
org.apache.activemq	activemq-client	5.15.0
org.apache.activemq	activemq-console	5.15.0
org.apache.activemq	activemq-http	5.15.0
org.apache.activemq	activemq-jaas	5.15.0
org.apache.activemq	activemq-jdbc-store	5.15.0
org.apache.activemq	activemq-jms-pool	5.15.0
org.apache.activemq	activemq-kahadb-store	5.15.0
org.apache.activemq	activemq-karaf	5.15.0
org.apache.activemq	activemq-leveldb-store	5.15.0
org.apache.activemq	activemq-log4j-appender	5.15.0
org.apache.activemq	activemq-mqtt	5.15.0
org.apache.activemq	activemq-openwire-generator	5.15.0
org.apache.activemq	activemq-openwire-legacy	5.15.0
org.apache.activemq	activemq-osgi	5.15.0
org.apache.activemq	activemq-partition	5.15.0
org.apache.activemq	activemq-pool	5.15.0
org.apache.activemq	activemq-ra	5.15.0
org.apache.activemq	activemq-run	5.15.0
org.apache.activemq	activemq-runtime-config	5.15.0
org.apache.activemq	activemq-shiro	5.15.0
org.apache.activemq	activemq-spring	5.15.0
org.apache.activemq	activemq-stomp	5.15.0
org.apache.activemq	activemq-web	5.15.0
org.apache.activemq	artemis-amqp-protocol	2.1.0
org.apache.activemq	artemis-commons	2.1.0
org.apache.activemq	artemis-core-client	2.1.0
org.apache.activemq	artemis-jms-client	2.1.0

Group ID	Artifact ID	Version
org.apache.activemq	artemis-jms-server	2.1.0
org.apache.activemq	artemis-journal	2.1.0
org.apache.activemq	artemis-native	2.1.0
org.apache.activemq	artemis-selector	2.1.0
org.apache.activemq	artemis-server	2.1.0
org.apache.activemq	artemis-service-extensions	2.1.0
org.apache.commons	commons-dbcp2	2.1.1
org.apache.commons	commons-lang3	3.6
org.apache.commons	commons-pool2	2.4.2
org.apache.derby	derby	10.13.1.1
org.apache.httpcomponents	httpasyncclient	4.1.3
org.apache.httpcomponents	httpclient	4.5.3
org.apache.httpcomponents	httpcore	4.4.6
org.apache.httpcomponents	httpcore-nio	4.4.6
org.apache.httpcomponents	httpmime	4.5.3
org.apache.logging.log4j	log4j-1.2-api	2.8.2
org.apache.logging.log4j	log4j-api	2.8.2
org.apache.logging.log4j	log4j-api-scala_2.10	2.8.2
org.apache.logging.log4j	log4j-api-scala_2.11	2.8.2
org.apache.logging.log4j	log4j-core	2.8.2
org.apache.logging.log4j	log4j-flume-ng	2.8.2
org.apache.logging.log4j	log4j-iostreams	2.8.2
org.apache.logging.log4j	log4j-jcl	2.8.2
org.apache.logging.log4j	log4j-jmx-gui	2.8.2
org.apache.logging.log4j	log4j-jul	2.8.2
org.apache.logging.log4j	log4j-liquibase	2.8.2
org.apache.logging.log4j	log4j-nosql	2.8.2
org.apache.logging.log4j	log4j-slf4j-impl	2.8.2
org.apache.logging.log4j	log4j-taglib	2.8.2
org.apache.logging.log4j	log4j-web	2.8.2

Group ID	Artifact ID	Version
org.apache.solr	solr-solrj	6.6.0
org.apache.tomcat	tomcat-jdbc	8.5.16
org.apache.tomcat	tomcat-jsp-api	8.5.16
org.apache.tomcat.embed	tomcat-embed-core	8.5.16
org.apache.tomcat.embed	tomcat-embed-el	8.5.16
org.apache.tomcat.embed	tomcat-embed-jasper	8.5.16
org.apache.tomcat.embed	tomcat-embed-websocket	8.5.16
org.aspectj	aspectjrt	1.8.10
org.aspectj	aspectjtools	1.8.10
org.aspectj	aspectjweaver	1.8.10
org.assertj	assertj-core	3.8.0
org.codehaus.btm	btm	2.1.4
org.codehaus.groovy	groovy	2.5.0-beta-1
org.codehaus.groovy	groovy-all	2.5.0-beta-1
org.codehaus.groovy	groovy-ant	2.5.0-beta-1
org.codehaus.groovy	groovy-bsf	2.5.0-beta-1
org.codehaus.groovy	groovy-console	2.5.0-beta-1
org.codehaus.groovy	groovy-docgenerator	2.5.0-beta-1
org.codehaus.groovy	groovy-groovydoc	2.5.0-beta-1
org.codehaus.groovy	groovy-groovysh	2.5.0-beta-1
org.codehaus.groovy	groovy-jmx	2.5.0-beta-1
org.codehaus.groovy	groovy-json	2.5.0-beta-1
org.codehaus.groovy	groovy-jsr223	2.5.0-beta-1
org.codehaus.groovy	groovy-nio	2.5.0-beta-1
org.codehaus.groovy	groovy-servlet	2.5.0-beta-1
org.codehaus.groovy	groovy-sql	2.5.0-beta-1
org.codehaus.groovy	groovy-swing	2.5.0-beta-1
org.codehaus.groovy	groovy-templates	2.5.0-beta-1
org.codehaus.groovy	groovy-test	2.5.0-beta-1
org.codehaus.groovy	groovy-testng	2.5.0-beta-1

Group ID	Artifact ID	Version
org.codehaus.groovy	groovy-xml	2.5.0-beta-1
org.codehaus.janino	janino	3.0.7
org.eclipse.jetty	apache-jsp	9.4.6.v20170531
org.eclipse.jetty	apache-jstl	9.4.6.v20170531
org.eclipse.jetty	jetty-annotations	9.4.6.v20170531
org.eclipse.jetty	jetty-client	9.4.6.v20170531
org.eclipse.jetty	jetty-continuation	9.4.6.v20170531
org.eclipse.jetty	jetty-deploy	9.4.6.v20170531
org.eclipse.jetty	jetty-http	9.4.6.v20170531
org.eclipse.jetty	jetty-http-spi	9.4.6.v20170531
org.eclipse.jetty	jetty-infinispan	9.4.6.v20170531
org.eclipse.jetty	jetty-io	9.4.6.v20170531
org.eclipse.jetty	jetty-jaas	9.4.6.v20170531
org.eclipse.jetty	jetty-jaspi	9.4.6.v20170531
org.eclipse.jetty	jetty-jmx	9.4.6.v20170531
org.eclipse.jetty	jetty-jndi	9.4.6.v20170531
org.eclipse.jetty	jetty-nosql	9.4.6.v20170531
org.eclipse.jetty	jetty-plus	9.4.6.v20170531
org.eclipse.jetty	jetty-proxy	9.4.6.v20170531
org.eclipse.jetty	jetty-quickstart	9.4.6.v20170531
org.eclipse.jetty	jetty-rewrite	9.4.6.v20170531
org.eclipse.jetty	jetty-runner	9.4.6.v20170531
org.eclipse.jetty	jetty-security	9.4.6.v20170531
org.eclipse.jetty	jetty-server	9.4.6.v20170531
org.eclipse.jetty	jetty-servlet	9.4.6.v20170531
org.eclipse.jetty	jetty-servlets	9.4.6.v20170531
org.eclipse.jetty	jetty-spring	9.4.6.v20170531
org.eclipse.jetty	jetty-start	9.4.6.v20170531
org.eclipse.jetty	jetty-util	9.4.6.v20170531
org.eclipse.jetty	jetty-util-ajax	9.4.6.v20170531

Group ID	Artifact ID	Version
org.eclipse.jetty	jetty-webapp	9.4.6.v20170531
org.eclipse.jetty	jetty-xml	9.4.6.v20170531
org.eclipse.jetty.http2	http2-client	9.4.6.v20170531
org.eclipse.jetty.http2	http2-common	9.4.6.v20170531
org.eclipse.jetty.http2	http2-hpack	9.4.6.v20170531
org.eclipse.jetty.http2	http2-server	9.4.6.v20170531
org.eclipse.jetty.orbit	javax.servlet.jsp	2.2.0.v201112011158
org.eclipse.jetty.websocket	javax-websocket-server-impl	9.4.6.v20170531
org.eclipse.jetty.websocket	websocket-client	9.4.6.v20170531
org.eclipse.jetty.websocket	websocket-server	9.4.6.v20170531
org.ehcache	ehcache	3.3.1
org.ehcache	ehcache-clustered	3.3.1
org.ehcache	ehcache-transactions	3.3.1
org.elasticsearch	elasticsearch	5.5.0
org.elasticsearch.client	transport	5.5.0
org.elasticsearch.plugin	transport-netty4-client	5.5.0
org.firebirdsql.jdbc	jaybird-jdk17	3.0.1
org.firebirdsql.jdbc	jaybird-jdk18	3.0.1
org.flywaydb	flyway-core	4.2.0
org.freemarker	freemarker	2.3.26-incubating
org.glassfish	javax.el	3.0.0
org.glassfish.jersey.bundles	jersey-guava	2.25.1
org.glassfish.jersey.containers	jersey-container-servlet	2.25.1
org.glassfish.jersey.containers	jersey-container-servlet-core	2.25.1
org.glassfish.jersey.core	jersey-client	2.25.1
org.glassfish.jersey.core	jersey-common	2.25.1
org.glassfish.jersey.core	jersey-server	2.25.1
org.glassfish.jersey.ext	jersey-bean-validation	2.25.1

Group ID	Artifact ID	Version
org.glassfish.jersey.ext	jersey-entity-filtering	2.25.1
org.glassfish.jersey.ext	jersey-spring3	2.25.1
org.glassfish.jersey.media	jersey-media-jaxb	2.25.1
org.glassfish.jersey.media	jersey-media-json-jackson	2.25.1
org.glassfish.jersey.media	jersey-media-multipart	2.25.1
org.hamcrest	hamcrest-core	1.3
org.hamcrest	hamcrest-library	1.3
org.hibernate	hibernate-core	5.2.10.Final
org.hibernate	hibernate-ehcache	5.2.10.Final
org.hibernate	hibernate-entitymanager	5.2.10.Final
org.hibernate	hibernate-envers	5.2.10.Final
org.hibernate	hibernate-java8	5.2.10.Final
org.hibernate	hibernate-jpamodelgen	5.2.10.Final
org.hibernate	hibernate-validator	5.4.1.Final
org.hibernate	hibernate-validator-annotation-processor	5.4.1.Final
org.hsqldb	hsqldb	2.4.0
org.infinispan	infinispan-jcache	9.1.0.Final
org.infinispan	infinispan-spring4-common	9.1.0.Final
org.infinispan	infinispan-spring4-embedded	9.1.0.Final
org.influxdb	influxdb-java	2.7
org.javassist	javassist	3.21.0-GA
org.jboss	jboss-transaction-spi	7.6.0.Final
org.jboss.logging	jboss-logging	3.3.1.Final
org.jboss.narayana.jta	jdbc	5.6.3.Final
org.jboss.narayana.jta	jms	5.6.3.Final
org.jboss.narayana.jta	jta	5.6.3.Final
org.jboss.narayana.jts	narayana-jts-integration	5.6.3.Final

Group ID	Artifact ID	Version
org.jdom	jdom2	2.0.6
org.jolokia	jolokia-core	1.3.7
org.jooq	jooq	3.9.4
org.jooq	jooq-codegen	3.9.4
org.jooq	jooq-meta	3.9.4
org.liquibase	liquibase-core	3.5.3
org.mariadb.jdbc	mariadb-java-client	2.0.3
org.mockito	mockito-core	2.8.47
org.mongodb	bson	3.4.2
org.mongodb	mongodb-driver	3.4.2
org.mongodb	mongodb-driver-async	3.4.2
org.mongodb	mongodb-driver-core	3.4.2
org.mongodb	mongodb-driver-reactivestreams	1.5.0
org.mongodb	mongo-java-driver	3.4.2
org.mortbay.jasper	apache-el	8.5.9.1
org.neo4j	neo4j-ogm-api	3.0.0-RC1
org.neo4j	neo4j-ogm-bolt-driver	3.0.0-RC1
org.neo4j	neo4j-ogm-core	3.0.0-RC1
org.neo4j	neo4j-ogm-http-driver	3.0.0-RC1
org.postgresql	postgresql	42.1.3
org.projectlombok	lombok	1.16.18
org.quartz-scheduler	quartz	2.3.0
org.reactivestreams	reactive-streams	1.0.0
org.seleniumhq.selenium	htmlunit-driver	2.27
org.seleniumhq.selenium	selenium-api	3.4.0
org.seleniumhq.selenium	selenium-chrome-driver	3.4.0
org.seleniumhq.selenium	selenium-firefox-driver	3.4.0
org.seleniumhq.selenium	selenium-ie-driver	3.4.0
org.seleniumhq.selenium	selenium-java	3.4.0
org.seleniumhq.selenium	selenium-remote-driver	3.4.0

Group ID	Artifact ID	Version
org.seleniumhq.selenium	selenium-safari-driver	3.4.0
org.seleniumhq.selenium	selenium-support	3.4.0
org.skyscreamer	jsonassert	1.5.0
org.slf4j	jcl-over-slf4j	1.7.25
org.slf4j	jul-to-slf4j	1.7.25
org.slf4j	log4j-over-slf4j	1.7.25
org.slf4j	slf4j-api	1.7.25
org.slf4j	slf4j-jdk14	1.7.25
org.slf4j	slf4j-log4j12	1.7.25
org.slf4j	slf4j-simple	1.7.25
org.springframework	spring-aop	5.0.0.RC3
org.springframework	spring-aspects	5.0.0.RC3
org.springframework	spring-beans	5.0.0.RC3
org.springframework	spring-context	5.0.0.RC3
org.springframework	spring-context-indexer	5.0.0.RC3
org.springframework	spring-context-support	5.0.0.RC3
org.springframework	spring-core	5.0.0.RC3
org.springframework	spring-expression	5.0.0.RC3
org.springframework	spring-instrument	5.0.0.RC3
org.springframework	spring-jcl	5.0.0.RC3
org.springframework	spring-jdbc	5.0.0.RC3
org.springframework	spring-jms	5.0.0.RC3
org.springframework	spring-messaging	5.0.0.RC3
org.springframework	spring-orm	5.0.0.RC3
org.springframework	spring-oxm	5.0.0.RC3
org.springframework	spring-test	5.0.0.RC3
org.springframework	spring-tx	5.0.0.RC3
org.springframework	spring-web	5.0.0.RC3
org.springframework	spring-webflux	5.0.0.RC3
org.springframework	spring-webmvc	5.0.0.RC3

Group ID	Artifact ID	Version
org.springframework	spring-websocket	5.0.0.RC3
org.springframework.amqp	spring-amqp	2.0.0.M5
org.springframework.amqp	spring-rabbit	2.0.0.M5
org.springframework.batch	spring-batch-core	4.0.0.M3
org.springframework.batch	spring-batch- infrastructure	4.0.0.M3
org.springframework.batch	spring-batch- integration	4.0.0.M3
org.springframework.batch	spring-batch-test	4.0.0.M3
org.springframework.boot	spring-boot	2.0.0.M3
org.springframework.boot	spring-boot-actuator	2.0.0.M3
org.springframework.boot	spring-boot-actuator- docs	2.0.0.M3
org.springframework.boot	spring-boot- autoconfigure	2.0.0.M3
org.springframework.boot	spring-boot- autoconfigure-processor	2.0.0.M3
org.springframework.boot	spring-boot- configuration-metadata	2.0.0.M3
org.springframework.boot	spring-boot- configuration-processor	2.0.0.M3
org.springframework.boot	spring-boot-devtools	2.0.0.M3
org.springframework.boot	spring-boot-loader	2.0.0.M3
org.springframework.boot	spring-boot-loader- tools	2.0.0.M3
org.springframework.boot	spring-boot-starter	2.0.0.M3
org.springframework.boot	spring-boot-starter- activemq	2.0.0.M3
org.springframework.boot	spring-boot-starter- actuator	2.0.0.M3
org.springframework.boot	spring-boot-starter- amqp	2.0.0.M3
org.springframework.boot	spring-boot-starter-aop	2.0.0.M3
org.springframework.boot	spring-boot-starter- artemis	2.0.0.M3

Group ID	Artifact ID	Version
org.springframework.boot	spring-boot-starter-batch	2.0.0.M3
org.springframework.boot	spring-boot-starter-cache	2.0.0.M3
org.springframework.boot	spring-boot-starter-cloud-connectors	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-cassandra	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-cassandra-reactive	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-couchbase	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-elasticsearch	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-jpa	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-ldap	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-mongodb	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-mongodb-reactive	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-neo4j	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-redis	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-redis-reactive	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-rest	2.0.0.M3
org.springframework.boot	spring-boot-starter-data-solr	2.0.0.M3
org.springframework.boot	spring-boot-starter-freemarker	2.0.0.M3
org.springframework.boot	spring-boot-starter-groovy-templates	2.0.0.M3

Group ID	Artifact ID	Version
org.springframework.boot	spring-boot-starter-hateoas	2.0.0.M3
org.springframework.boot	spring-boot-starter-integration	2.0.0.M3
org.springframework.boot	spring-boot-starter-jdbc	2.0.0.M3
org.springframework.boot	spring-boot-starter-jersey	2.0.0.M3
org.springframework.boot	spring-boot-starter-jetty	2.0.0.M3
org.springframework.boot	spring-boot-starter-jooq	2.0.0.M3
org.springframework.boot	spring-boot-starter-json	2.0.0.M3
org.springframework.boot	spring-boot-starter-jta-atomikos	2.0.0.M3
org.springframework.boot	spring-boot-starter-jta-bitronix	2.0.0.M3
org.springframework.boot	spring-boot-starter-jta-narayana	2.0.0.M3
org.springframework.boot	spring-boot-starter-log4j2	2.0.0.M3
org.springframework.boot	spring-boot-starter-logging	2.0.0.M3
org.springframework.boot	spring-boot-starter-mail	2.0.0.M3
org.springframework.boot	spring-boot-starter-mobile	2.0.0.M3
org.springframework.boot	spring-boot-starter-mustache	2.0.0.M3
org.springframework.boot	spring-boot-starter-quartz	2.0.0.M3
org.springframework.boot	spring-boot-starter-reactor-netty	2.0.0.M3
org.springframework.boot	spring-boot-starter-security	2.0.0.M3

Group ID	Artifact ID	Version
org.springframework.boot	spring-boot-starter-social-facebook	2.0.0.M3
org.springframework.boot	spring-boot-starter-social-linkedin	2.0.0.M3
org.springframework.boot	spring-boot-starter-social-twitter	2.0.0.M3
org.springframework.boot	spring-boot-starter-test	2.0.0.M3
org.springframework.boot	spring-boot-starter-thymeleaf	2.0.0.M3
org.springframework.boot	spring-boot-starter-tomcat	2.0.0.M3
org.springframework.boot	spring-boot-starter-undertow	2.0.0.M3
org.springframework.boot	spring-boot-starter-validation	2.0.0.M3
org.springframework.boot	spring-boot-starter-web	2.0.0.M3
org.springframework.boot	spring-boot-starter-webflux	2.0.0.M3
org.springframework.boot	spring-boot-starter-web-services	2.0.0.M3
org.springframework.boot	spring-boot-starter-websocket	2.0.0.M3
org.springframework.boot	spring-boot-test	2.0.0.M3
org.springframework.boot	spring-boot-test-autoconfigure	2.0.0.M3
org.springframework.cloud	spring-cloud-cloudfoundry-connector	2.0.0.M1
org.springframework.cloud	spring-cloud-core	2.0.0.M1
org.springframework.cloud	spring-cloud-heroku-connector	2.0.0.M1
org.springframework.cloud	spring-cloud-localconfig-connector	2.0.0.M1
org.springframework.cloud	spring-cloud-spring-service-connector	2.0.0.M1
org.springframework.data	spring-data-cassandra	2.0.0.RC2

Group ID	Artifact ID	Version
org.springframework.data	spring-data-commons	2.0.0.RC2
org.springframework.data	spring-data-couchbase	3.0.0.RC2
org.springframework.data	spring-data- elasticsearch	3.0.0.RC2
org.springframework.data	spring-data-envers	2.0.0.RC2
org.springframework.data	spring-data-gemfire	2.0.0.RC2
org.springframework.data	spring-data-geode	2.0.0.RC2
org.springframework.data	spring-data-jpa	2.0.0.RC2
org.springframework.data	spring-data-keyvalue	2.0.0.RC2
org.springframework.data	spring-data-ldap	2.0.0.RC2
org.springframework.data	spring-data-mongodb	2.0.0.RC2
org.springframework.data	spring-data-mongodb- cross-store	2.0.0.RC2
org.springframework.data	spring-data-neo4j	5.0.0.RC2
org.springframework.data	spring-data-redis	2.0.0.RC2
org.springframework.data	spring-data-rest-core	3.0.0.RC2
org.springframework.data	spring-data-rest-hal- browser	3.0.0.RC2
org.springframework.data	spring-data-rest-webmvc	3.0.0.RC2
org.springframework.data	spring-data-solr	3.0.0.RC2
org.springframework.hateoas	spring-hateoas	0.23.0.RELEASE
org.springframework.integration	spring-integration-amqp	5.0.0.M6
org.springframework.integration	spring-integration-core	5.0.0.M6
org.springframework.integration	spring-integration- event	5.0.0.M6
org.springframework.integration	spring-integration-feed	5.0.0.M6
org.springframework.integration	spring-integration-file	5.0.0.M6
org.springframework.integration	spring-integration-ftp	5.0.0.M6
org.springframework.integration	spring-integration- gemfire	5.0.0.M6
org.springframework.integration	spring-integration- groovy	5.0.0.M6

Group ID	Artifact ID	Version
org.springframework.integration	spring-integration-http	5.0.0.M6
org.springframework.integration	spring-integration-ip	5.0.0.M6
org.springframework.integration	spring-integration-jdbc	5.0.0.M6
org.springframework.integration	spring-integration-jms	5.0.0.M6
org.springframework.integration	spring-integration-jmx	5.0.0.M6
org.springframework.integration	spring-integration-jpa	5.0.0.M6
org.springframework.integration	spring-integration-mail	5.0.0.M6
org.springframework.integration	spring-integration-mongodb	5.0.0.M6
org.springframework.integration	spring-integration-mqtt	5.0.0.M6
org.springframework.integration	spring-integration-redis	5.0.0.M6
org.springframework.integration	spring-integration-rmi	5.0.0.M6
org.springframework.integration	spring-integration-scripting	5.0.0.M6
org.springframework.integration	spring-integration-security	5.0.0.M6
org.springframework.integration	spring-integration-sftp	5.0.0.M6
org.springframework.integration	spring-integration-stomp	5.0.0.M6
org.springframework.integration	spring-integration-stream	5.0.0.M6
org.springframework.integration	spring-integration-syslog	5.0.0.M6
org.springframework.integration	spring-integration-test	5.0.0.M6
org.springframework.integration	spring-integration-test-support	5.0.0.M6
org.springframework.integration	spring-integration-twitter	5.0.0.M6
org.springframework.integration	spring-integration-websocket	5.0.0.M6
org.springframework.integration	spring-integration-ws	5.0.0.M6
org.springframework.integration	spring-integration-xml	5.0.0.M6
org.springframework.integration	spring-integration-xmpp	5.0.0.M6

Group ID	Artifact ID	Version
org.springframework.integration	spring-integration-zookeeper	5.0.0.M6
org.springframework.kafka	spring-kafka	2.0.0.M3
org.springframework.kafka	spring-kafka-test	2.0.0.M3
org.springframework.ldap	spring-ldap-core	2.3.1.RELEASE
org.springframework.ldap	spring-ldap-core-tiger	2.3.1.RELEASE
org.springframework.ldap	spring-ldap-ldif-batch	2.3.1.RELEASE
org.springframework.ldap	spring-ldap-ldif-core	2.3.1.RELEASE
org.springframework.ldap	spring-ldap-odm	2.3.1.RELEASE
org.springframework.ldap	spring-ldap-test	2.3.1.RELEASE
org.springframework.mobile	spring-mobile-device	2.0.0.M1
org.springframework.plugin	spring-plugin-core	1.2.0.RELEASE
org.springframework.plugin	spring-plugin-metadata	1.2.0.RELEASE
org.springframework.restdocs	spring-restdocs-asciidoctor	1.2.1.RELEASE
org.springframework.restdocs	spring-restdocs-core	1.2.1.RELEASE
org.springframework.restdocs	spring-restdocs-mockmvc	1.2.1.RELEASE
org.springframework.restdocs	spring-restdocs-restassured	1.2.1.RELEASE
org.springframework.retry	spring-retry	1.2.1.RELEASE
org.springframework.security	spring-security-acl	5.0.0.M3
org.springframework.security	spring-security-aspects	5.0.0.M3
org.springframework.security	spring-security-cas	5.0.0.M3
org.springframework.security	spring-security-config	5.0.0.M3
org.springframework.security	spring-security-core	5.0.0.M3
org.springframework.security	spring-security-crypto	5.0.0.M3
org.springframework.security	spring-security-data	5.0.0.M3
org.springframework.security	spring-security-jwt	1.0.8.RELEASE
org.springframework.security	spring-security-jwt-jose	5.0.0.M3
org.springframework.security	spring-security-ldap	5.0.0.M3

Group ID	Artifact ID	Version
org.springframework.security	spring-security-messaging	5.0.0.M3
org.springframework.security	spring-security-oauth2-client	5.0.0.M3
org.springframework.security	spring-security-oauth2-core	5.0.0.M3
org.springframework.security	spring-security-openid	5.0.0.M3
org.springframework.security	spring-security-remoting	5.0.0.M3
org.springframework.security	spring-security-taglibs	5.0.0.M3
org.springframework.security	spring-security-test	5.0.0.M3
org.springframework.security	spring-security-web	5.0.0.M3
org.springframework.security	spring-security-webflux	5.0.0.M3
org.springframework.security	spring-security-oauth	2.2.0.RC1
org.springframework.security	spring-security-oauth2	2.2.0.RC1
org.springframework.session	spring-session-core	2.0.0.M3
org.springframework.session	spring-session-data-redis	2.0.0.M3
org.springframework.session	spring-session-hazelcast	2.0.0.M3
org.springframework.session	spring-session-jdbc	2.0.0.M3
org.springframework.social	spring-social-config	2.0.0.M3
org.springframework.social	spring-social-core	2.0.0.M3
org.springframework.social	spring-social-facebook	3.0.0.M2
org.springframework.social	spring-social-facebook-web	3.0.0.M2
org.springframework.social	spring-social-linkedin	2.0.0.M2
org.springframework.social	spring-social-security	2.0.0.M3
org.springframework.social	spring-social-twitter	2.0.0.M2
org.springframework.social	spring-social-web	2.0.0.M3
org.springframework.ws	spring-ws-core	2.4.0.RELEASE
org.springframework.ws	spring-ws-security	2.4.0.RELEASE
org.springframework.ws	spring-ws-support	2.4.0.RELEASE

Group ID	Artifact ID	Version
org.springframework.ws	spring-ws-test	2.4.0.RELEASE
org.thymeleaf	thymeleaf	3.0.7.RELEASE
org.thymeleaf	thymeleaf-spring5	3.0.7.RC1
org.thymeleaf.extras	thymeleaf-extras-java8time	3.0.0.RELEASE
org.thymeleaf.extras	thymeleaf-extras-springsecurity4	3.0.2.RELEASE
org.webjars	hal-browser	3325375
org.webjars	webjars-locator	0.32-1
org.xerial	sqlite-jdbc	3.19.3
org.yaml	snakeyaml	1.18
redis.clients	jedis	2.9.0
wsdl4j	wsdl4j	1.6.3
xml-apis	xml-apis	1.4.01