
Dropwizard Documentation

Release 4.0.1-SNAPSHOT

Coda Hale

Jun 02, 2023

GENERAL

1 Security	3
2 Frequently Asked Questions	5
3 Getting Started	7
4 Dropwizard Core	21
5 Dropwizard Dependency Injection	59
6 Dropwizard Client	63
7 Dropwizard JDBC	69
8 Dropwizard Migrations	73
9 Dropwizard Hibernate	81
10 Dropwizard Authentication	85
11 Dropwizard Forms	93
12 Dropwizard Validation	95
13 Dropwizard Views	105
14 Dropwizard & Scala	111
15 Testing Dropwizard	113
16 Upgrade Notes	125
17 Dropwizard Example, Step by Step	143
18 Dropwizard Configuration Reference	145
19 Dropwizard Internals	181
20 Contributors	185
21 Sponsors	197
22 Other Versions	199

Dropwizard pulls together **stable, mature** libraries from the Java ecosystem into a **simple, light-weight** package that lets you focus on *getting things done*.

Dropwizard has *out-of-the-box* support for sophisticated **configuration, application metrics, logging, operational tools**, and much more, allowing you and your team to ship a *production-quality* web service in the shortest time possible.

SECURITY

No known issues exist

FREQUENTLY ASKED QUESTIONS

What's a Dropwizard?

A character in a [K.C. Green web comic](#).

How is Dropwizard licensed?

It's licensed under the [Apache License v2](#).

How can I commit to Dropwizard?

Go to the [GitHub project](#), fork it, and submit a pull request. We prefer small, single-purpose pull requests over large, multi-purpose ones. We reserve the right to turn down any proposed changes, but in general we're delighted when people want to make our projects better!

GETTING STARTED

Getting Started will guide you through the process of creating a simple Dropwizard Project: Hello World. Along the way, we'll explain the various underlying libraries and their roles, important concepts in Dropwizard, and suggest some organizational techniques to help you as your project grows. (Or you can just skip to the fun part.)

3.1 Overview

Dropwizard straddles the line between being a library and a framework. Its goal is to provide performant, reliable implementations of everything a production-ready web application needs. Because this functionality is extracted into a reusable library, your application remains lean and focused, reducing both time-to-market and maintenance burdens.

3.1.1 Jetty for HTTP

Because you can't be a web application without HTTP, Dropwizard uses the [Jetty](#) HTTP library to embed an incredibly tuned HTTP server directly into your project. Instead of handing your application off to a complicated application server, Dropwizard projects have a `main` method which spins up an HTTP server. Running your application as a simple process eliminates a number of unsavory aspects of Java in production (no PermGen issues, no application server configuration and maintenance, no arcane deployment tools, no class loader troubles, no hidden application logs, no trying to tune a single garbage collector to work with multiple application workloads) and allows you to use all of the existing Unix process management tools instead.

3.1.2 Jersey for REST

For building RESTful web applications, we've found nothing beats [Jersey](#) (the [JAX-RS](#) reference implementation) in terms of features or performance. It allows you to write clean, testable classes which gracefully map HTTP requests to simple Java objects. It supports streaming output, matrix URI parameters, conditional GET requests, and much, much more.

3.1.3 Jackson for JSON

In terms of data formats, JSON has become the web's *lingua franca*, and [Jackson](#) is the king of JSON on the JVM. In addition to being lightning fast, it has a sophisticated object mapper, allowing you to export your domain models directly.

3.1.4 Metrics for metrics

The [Metrics](#) library rounds things out, providing you with unparalleled insight into your code's behavior in your production environment.

3.1.5 And Friends

In addition to [Jetty](#), [Jersey](#), and [Jackson](#), Dropwizard also includes a number of libraries to help you ship more quickly and with fewer regrets.

- [Logback](#) and [slf4j](#) for performant and flexible logging.
- [Hibernate Validator](#), the JSR 349 reference implementation, provides an easy, declarative framework for validating user input and generating helpful and i18n-friendly error messages.
- The [Apache HttpClient](#) and [Jersey](#) client libraries allow for both low- and high-level interaction with other web services.
- [JDBI](#) is the most straightforward way to use a relational database with Java.
- [Liquibase](#) is a great way to keep your database schema in check throughout your development and release cycles, applying high-level database refactorings instead of one-off DDL scripts.
- [Freemarker](#) and [Mustache](#) are simple templating systems for more user-facing applications.

Now that you've gotten the lay of the land, let's dig in!

3.2 Setting Up Using Maven

We recommend you use [Maven](#) for new Dropwizard applications. If you're a big [Ant](#) / [Ivy](#), [Buildr](#), [Gradle](#), [SBT](#), [Leiningen](#), or [Gant](#) fan, that's cool, but we use Maven, and we'll be using Maven as we go through this example application. If you have any questions about how Maven works, [Maven: The Complete Reference](#) should have what you're looking for.

You have three alternatives from here:

1. Create a project using [dropwizard-archetype](#):

```
mvn archetype:generate -DarchetypeGroupId=io.dropwizard.archetypes -  
↪DarchetypeArtifactId=java-simple -DarchetypeVersion=[REPLACE WITH A VALID↪  
↪DROPWIZARD VERSION]
```

2. Look at the [dropwizard-example](#)
3. Follow the tutorial below to see how you can include it in your existing project

3.2.1 Tutorial

First, add the `dropwizard-bom` bill of materials (BOM) into the `dependencyManagement` section of your POM with the current version of Dropwizard (which is 4.0.1-SNAPSHOT):

Listing 1: Source file

```
<dependencyManagement>
  <dependencies>
    <dependency>
      <groupId>io.dropwizard</groupId>
      <artifactId>dropwizard-bom</artifactId>
      <version>4.0.1-SNAPSHOT</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

Add the `dropwizard-core` library as a dependency:

Listing 2: Source file

```
<dependencies>
  <dependency>
    <groupId>io.dropwizard</groupId>
    <artifactId>dropwizard-core</artifactId>
  </dependency>
</dependencies>
```

Alright, that's enough XML. We've got a Maven project set up now, and it's time to start writing real code.

3.3 Creating A Configuration Class

Each Dropwizard application has its own subclass of the `Configuration` class which specifies environment-specific parameters. These parameters are specified in a **YAML** configuration file which is deserialized to an instance of your application's configuration class and validated.

The application we'll be building is a high-performance Hello World service, and one of our requirements is that we need to be able to vary how it says hello from environment to environment. We'll need to specify at least two things to begin with: a template for saying hello and a default name to use in case the user doesn't specify their name.

Here's what our configuration class will look like, full [example conf](#) here:

Listing 3: Source file

```
package com.example.helloworld;

import com.fasterxml.jackson.annotation.JsonProperty;
import io.dropwizard.core.Configuration;
import jakarta.validation.constraints.NotEmpty;

public class HelloWorldConfiguration extends Configuration {
    @NotEmpty
```

(continues on next page)

(continued from previous page)

```
private String template;

@NotEmpty
private String defaultName = "Stranger";

@JsonProperty
public String getTemplate() {
    return template;
}

@JsonProperty
public void setTemplate(String template) {
    this.template = template;
}

@JsonProperty
public String getDefaultName() {
    return defaultName;
}

@JsonProperty
public void setDefaultName(String name) {
    this.defaultName = name;
}
}
```

There's a lot going on here, so let's unpack a bit of it.

When this class is deserialized from the YAML file, it will pull two root-level fields from the YAML object: `template`, the template for our Hello World saying, and `defaultName`, the default name to use. Both `template` and `defaultName` are annotated with `@NotEmpty`, so if the YAML configuration file has blank values for either or is missing `template` entirely an informative exception will be thrown, and your application won't start.

Both the getters and setters for `template` and `defaultName` are annotated with `@JsonProperty`, which allows Jackson to both deserialize the properties from a YAML file but also to serialize it.

Note: The mapping from YAML to your application's `Configuration` instance is done by `Jackson`. This means your `Configuration` class can use all of Jackson's `object-mapping annotations`. The validation of `@NotEmpty` is handled by Hibernate Validator, which has a `wide range of built-in constraints` for you to use.

Our YAML file will then look like the below, full `example.yml` here:

Listing 4: Source file

```
---
template: Hello, %s!
defaultName: Stranger
```

Dropwizard has *many* more configuration parameters than that, but they all have sane defaults so you can keep your configuration files small and focused.

So save that YAML file in the directory you plan to run the fat jar from (see below) as `hello-world.yml`, because we'll be getting up and running pretty soon, and we'll need it. Next up, we're creating our application class!

3.4 Creating An Application Class

Combined with your project's Configuration subclass, its Application subclass forms the core of your Dropwizard application. The Application class pulls together the various bundles and commands which provide basic functionality. (More on that later.) For now, though, our HelloWorldApplication looks like this:

Listing 5: Source file

```
package com.example.helloworld;

import io.dropwizard.core.Application;
import io.dropwizard.core.setup.Bootstrap;
import io.dropwizard.core.setup.Environment;

public class HelloWorldApplicationStart extends Application<HelloWorldConfiguration> {
    public static void main(String[] args) throws Exception {
        new HelloWorldApplicationStart().run(args);
    }

    @Override
    public String getName() {
        return "hello-world";
    }

    @Override
    public void initialize(Bootstrap<HelloWorldConfiguration> bootstrap) {
        // nothing to do yet
    }

    @Override
    public void run(HelloWorldConfiguration configuration, Environment environment) {
        // nothing to do yet
    }
}
```

As you can see, HelloWorldApplication is parameterized with the application's configuration type, HelloWorldConfiguration. An initialize method is used to configure aspects of the application required before the application is run, like bundles, configuration source providers, etc. Also, we've added a static main method, which will be our application's entry point. Right now, we don't have any functionality implemented, so our run method is a little boring. Let's fix that!

3.5 Creating A Representation Class

Before we can get into the nuts-and-bolts of our Hello World application, we need to stop and think about our API. Luckily, our application needs to conform to an industry standard, RFC 1149, which specifies the following JSON representation of a Hello World saying:

```
{
  "id": 1,
  "content": "Hi!"
}
```

The `id` field is a unique identifier for the saying, and `content` is the textual representation of the saying. (Thankfully, this is a fairly straight-forward industry standard.)

To model this representation, we'll create a representation class:

Listing 6: Source file

```
package com.example.helloworld.api;

import com.fasterxml.jackson.annotation.JsonProperty;

public class Saying {
    private long id;
    private String content;

    public Saying() {
        // Jackson deserialization
    }

    public Saying(long id, String content) {
        this.id = id;
        this.content = content;
    }

    @JsonProperty
    public long getId() {
        return id;
    }

    @JsonProperty
    public String getContent() {
        return content;
    }
}
```

This is a pretty simple POJO, but there are a few things worth noting here.

First, it's immutable. This makes `Saying` instances *very* easy to reason about in multi-threaded environments as well as single-threaded environments. Second, it uses the JavaBeans standard for the `id` and `content` properties. This allows [Jackson](#) to serialize it to the JSON we need. The Jackson object mapping code will populate the `id` field of the JSON object with the return value of `#getId()`, likewise with `content` and `#getContent()`.

Note: The JSON serialization here is done by Jackson, which supports far more than simple JavaBean objects like this one. In addition to the sophisticated set of [annotations](#), you can even write your custom serializers and deserializers.

Now that we've got our representation class, it makes sense to start in on the resource it represents.

3.6 Creating A Resource Class

Jersey resources are the meat-and-potatoes of a Dropwizard application. Each resource class is associated with a URI template. For our application, we need a resource which returns new `Saying` instances from the URI `/hello-world`, so our resource class looks like this:

Listing 7: Source file

```
package com.example.helloworld.resources;

import com.codahale.metrics.annotation.Timed;
import com.example.helloworld.api.Saying;
import jakarta.ws.rs.GET;
import jakarta.ws.rs.Path;
import jakarta.ws.rs.Produces;
import jakarta.ws.rs.QueryParam;
import jakarta.ws.rs.core.MediaType;

import java.util.Optional;
import java.util.concurrent.atomic.AtomicLong;

@Path("/hello-world")
@Produces(MediaType.APPLICATION_JSON)
public class HelloWorldResource {
    private final String template;
    private final String defaultName;
    private final AtomicLong counter;

    public HelloWorldResource(String template, String defaultName) {
        this.template = template;
        this.defaultName = defaultName;
        this.counter = new AtomicLong();
    }

    @GET
    @Timed
    public Saying sayHello(@QueryParam("name") Optional<String> name) {
        final String value = String.format(template, name.orElse(defaultName));
        return new Saying(counter.incrementAndGet(), value);
    }
}
```

Finally, we're in the thick of it! Let's start from the top and work our way down.

`HelloWorldResource` has two annotations: `@Path` and `@Produces`. `@Path("/hello-world")` tells Jersey that this resource is accessible at the URI `/hello-world`, and `@Produces(MediaType.APPLICATION_JSON)` lets Jersey's content negotiation code know that this resource produces representations which are `application/json`.

`HelloWorldResource` takes two parameters for construction: the `template` it uses to produce the `saying` and the `defaultName` used when the user declines to tell us their name. An `AtomicLong` provides us with a cheap, thread-safe way of generating unique(ish) IDs.

Warning: Resource classes are used by multiple threads concurrently. In general, we recommend that resources be stateless/immutable, but it's important to keep the context in mind.

`#sayHello(Optional<String>)` is the meat of this class, and it's a fairly simple method. The `@QueryParam("name")` annotation tells Jersey to map the `name` parameter from the query string to the `name` parameter in the method. If the client sends a request to `/hello-world?name=Dougie`, `sayHello` will be called with `Optional.of("Dougie")`; if there is no `name` parameter in the query string, `sayHello` will be called with `Optional.empty()`. (Support for `Optional` is a little extra sauce that Dropwizard adds to Jersey's existing functionality.)

Note: If the client sends a request to `/hello-world?name=`, `sayHello` will be called with `Optional.of("")`. This may seem odd at first, but this follows the standards (an application may have different behavior depending on if a parameter is empty vs nonexistent). You can swap `Optional<String>` parameter with `NonEmptyStringParam` if you want `/hello-world?name=` to return "Hello, Stranger!" For more information on resource parameters see [the documentation](#)

Inside the `sayHello` method, we increment the counter, format the template using `String.format(String, Object...)`, and return a new `Saying` instance.

Because `sayHello` is annotated with `@Timed`, Dropwizard automatically records the duration and rate of its invocations as a Metrics Timer.

Once `sayHello` has returned, Jersey takes the `Saying` instance and looks for a provider class which can write `Saying` instances as `application/json`. Dropwizard has one such provider built in which allows for producing and consuming Java objects as JSON objects. The provider writes out the JSON and the client receives a `200 OK` response with a content type of `application/json`.

3.6.1 Registering A Resource

Before that will actually work, though, we need to go back to `HelloWorldApplication` and add this new resource class. In its `run` method we can read the template and default name from the `HelloWorldConfiguration` instance, create a new `HelloWorldResource` instance, and then add it to the application's Jersey environment:

Listing 8: Source file

```
HelloWorldResource resource = new HelloWorldResource(
    configuration.getTemplate(),
    configuration.getDefaultName()
);
environment.jersey().register(resource);
```

When our application starts, we create a new instance of our resource class with the parameters from the configuration file and hand it off to the `Environment`, which acts like a registry of all the things your application can do.

Note: A Dropwizard application can contain *many* resource classes, each corresponding to its own URI pattern. Just add another `@Path`-annotated resource class and call `register` with an instance of the new class.

Before we go too far, we should add a health check for our application.

3.7 Creating A Health Check

Health checks give you a way of adding small tests to your application to allow you to verify that your application is functioning correctly in production. We **strongly** recommend that all of your applications have at least a minimal set of health checks.

Note: We recommend this so strongly, in fact, that Dropwizard will nag you should you neglect to add a health check to your project.

Since formatting strings is not likely to fail while an application is running (unlike, say, a database connection pool), we'll have to get a little creative here. We'll add a health check to make sure we can actually format the provided template:

Listing 9: Source file

```
package com.example.helloworld.health;

import com.codahale.metrics.health.HealthCheck;

public class TemplateHealthCheck extends HealthCheck {
    private final String template;

    public TemplateHealthCheck(String template) {
        this.template = template;
    }

    @Override
    protected Result check() throws Exception {
        final String saying = String.format(template, "TEST");
        if (!saying.contains("TEST")) {
            return Result.unhealthy("template doesn't include a name");
        }
        return Result.healthy();
    }
}
```

TemplateHealthCheck checks for two things: that the provided template is actually a well-formed format string, and that the template actually produces output with the given name.

If the string is not a well-formed format string (for example, someone accidentally put `Hello, %s%` in the configuration file), then `String.format(String, Object...)` will throw an `IllegalFormatException` and the health check will implicitly fail. If the rendered saying doesn't include the test string, the health check will explicitly fail by returning an unhealthy `Result`.

3.7.1 Adding A Health Check

As with most things in Dropwizard, we create a new instance with the appropriate parameters and add it to the Environment:

Listing 10: Source file

```
TemplateHealthCheck healthCheck = new TemplateHealthCheck(configuration.getTemplate());
environment.healthChecks().register("template", healthCheck);
```

Now we're almost ready to go!

3.8 Building Fat JARs

We recommend that you build your Dropwizard applications as “fat” JAR files — single .jar files which contain *all* of the .class files required to run your application. This allows you to build a single deployable artifact which you can promote from your staging environment to your QA environment to your production environment without worrying about differences in installed libraries. To start building our Hello World application as a fat JAR, we need to configure a Maven plugin called maven-shade. In the <build><plugins> section of your pom.xml file, add this:

Listing 11: Source file

```
<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-shade-plugin</artifactId>
  <version>3.2.4</version>
  <configuration>
    <createDependencyReducedPom>>true</createDependencyReducedPom>
    <filters>
      <filter>
        <artifact>*.*</artifact>
        <excludes>
          <exclude>META-INF/*.SF</exclude>
          <exclude>META-INF/*.DSA</exclude>
          <exclude>META-INF/*.RSA</exclude>
        </excludes>
      </filter>
    </filters>
  </configuration>
  <executions>
    <execution>
      <phase>package</phase>
      <goals>
        <goal>shade</goal>
      </goals>
      <configuration>
        <transformers>
          <transformer implementation="org.apache.maven.plugins.shade.resource.
↪ServicesResourceTransformer" />
          <transformer implementation="org.apache.maven.plugins.shade.resource.
↪ManifestResourceTransformer">
            <mainClass>com.example.helloworld.HelloWorldApplication</
```

(continues on next page)

(continued from previous page)

```

<mainClass>
    </transformer>
  </transformers>
</configuration>
</execution>
</executions>
</plugin>

```

This configures Maven to do a couple of things during its package phase:

- Produce a `pom.xml` file which doesn't include dependencies for the libraries whose contents are included in the fat JAR.
- Exclude all digital signatures from signed JARs. If you don't, then Java considers the signature invalid and won't load or run your JAR file.
- Collate the various `META-INF/services` entries in the JARs instead of overwriting them. (Neither Dropwizard nor Jersey works without those.)
- Set `com.example.helloworld.HelloWorldApplication` as the JAR's `MainClass`. This will allow you to run the JAR using `java -jar`.

Warning: If your application has a dependency which *must* be signed (e.g., a [JCA/JCE](#) provider or other trusted library), you have to add an `exclusion` to the `maven-shade-plugin` configuration for that library and include that JAR in the classpath.

Warning: Since Dropwizard is using the Java `ServiceLoader` functionality to register and load extensions, the `minimizeJar` option of the `maven-shade-plugin` will lead to non-working application JARs.

3.8.1 Versioning Your JARs

Dropwizard can also use the project version if it's embedded in the JAR's manifest as the `Implementation-Version`. To embed this information using Maven, add the following to the `<build><plugins>` section of your `pom.xml` file:

Listing 12: Source file

```

<plugin>
  <groupId>org.apache.maven.plugins</groupId>
  <artifactId>maven-jar-plugin</artifactId>
  <version>3.2.0</version>
  <configuration>
    <archive>
      <manifest>
        <addDefaultImplementationEntries>true</addDefaultImplementationEntries>
      </manifest>
    </archive>
  </configuration>
</plugin>

```

This can be handy when trying to figure out what version of your application you have deployed on a machine.

Once you've got that configured, go into your project directory and run `mvn package` (or run the package goal from your IDE). You should see something like this:

```
[INFO] Including org.eclipse.jetty:jetty-util:jar:7.6.0.RC0 in the shaded jar.
[INFO] Including com.google.guava:guava:jar:10.0.1 in the shaded jar.
[INFO] Including com.google.code.findbugs:jsr305:jar:1.3.9 in the shaded jar.
[INFO] Including org.hibernate:hibernate-validator:jar:4.2.0.Final in the shaded jar.
[INFO] Including javax.validation:validation-api:jar:1.0.0.GA in the shaded jar.
[INFO] Including org.yaml:snakeyaml:jar:1.9 in the shaded jar.
[INFO] Replacing original artifact with shaded artifact.
[INFO] Replacing /Users/yourname/Projects/hello-world/target/hello-world-0.0.1-SNAPSHOT.
↪jar with /Users/yourname/Projects/hello-world/target/hello-world-0.0.1-SNAPSHOT-shaded.
↪jar
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----
[INFO] Total time: 8.415s
[INFO] Finished at: Fri Dec 02 16:26:42 PST 2011
[INFO] Final Memory: 11M/81M
[INFO] -----
```

Congratulations! You've built your first Dropwizard project! Now it's time to run it!

3.9 Running Your Application

Now that you've built a JAR file, it's time to run it.

In your project directory, run this:

```
java -jar target/hello-world-0.0.1-SNAPSHOT.jar
```

You should see something like the following:

```
usage: java -jar hello-world-0.0.1-SNAPSHOT.jar
       [-h] [-v] {server} ...

positional arguments:
  {server}              available commands

optional arguments:
  -h, --help            show this help message and exit
  -v, --version         show the service version and exit
```

Dropwizard takes the first command line argument and dispatches it to a matching command. In this case, the only command available is `server`, which runs your application as an HTTP server. The `server` command requires a configuration file, so let's go ahead and give it *the YAML file we previously saved*:

```
java -jar target/hello-world-0.0.1-SNAPSHOT.jar server hello-world.yml
```

You should see something like the following:

```
INFO [2011-12-03 00:38:32,927] io.dropwizard.core.cli.ServerCommand: Starting hello-
↪world
```

(continues on next page)

(continued from previous page)

```

INFO [2011-12-03 00:38:32,931] org.eclipse.jetty.server.Server: jetty-7.x.y-SNAPSHOT
INFO [2011-12-03 00:38:32,936] org.eclipse.jetty.server.handler.ContextHandler: started
↳o.e.j.s.ServletContextHandler{/,null}
INFO [2011-12-03 00:38:32,999] com.sun.jersey.server.impl.application.
↳WebApplicationImpl: Initiating Jersey application, version 'Jersey: 1.10 11/02/2011
↳03:53 PM'
INFO [2011-12-03 00:38:33,041] io.dropwizard.core.setup.Environment:

    GET    /hello-world (com.example.helloworld.resources.HelloWorldResource)

INFO [2011-12-03 00:38:33,215] org.eclipse.jetty.server.handler.ContextHandler: started
↳o.e.j.s.ServletContextHandler{/,null}
INFO [2011-12-03 00:38:33,235] org.eclipse.jetty.server.AbstractConnector: Started
↳BlockingChannelConnector@0.0.0.0:8080 STARTING
INFO [2011-12-03 00:38:33,238] org.eclipse.jetty.server.AbstractConnector: Started
↳SocketConnector@0.0.0.0:8081 STARTING

```

Your Dropwizard application is now listening on ports 8080 for application requests and 8081 for administration requests. If you press `^C`, the application will shut down gracefully, first closing the server socket, then waiting for in-flight requests to be processed, then shutting down the process itself.

However, while it's up, let's give it a whirl! [Click here to say hello!](#) [Click here to get even friendlier!](#)

So, we're generating sayings. Awesome. But that's not all your application can do. One of the main reasons for using Dropwizard is the out-of-the-box operational tools it provides, all of which can be found [on the admin port](#).

If you click through to the [metrics resource](#), you can see all of your application's metrics represented as a JSON object.

The [threads resource](#) allows you to quickly get a thread dump of all the threads running in that process.

Hint: When a Jetty worker thread is handling an incoming HTTP request, the thread name is set to the method and URI of the request. This can be *very* helpful when debugging a poorly-behaving request.

The [healthcheck resource](#) runs the *health check class we wrote*. You should see something like this:

```

* deadlocks: OK
* template: OK

```

`template` here is the result of your `TemplateHealthCheck`, which unsurprisingly passed. `deadlocks` is a built-in health check which looks for deadlocked JVM threads and prints out a listing if any are found.

3.10 Next Steps

Well, congratulations. You've got a Hello World application ready for production (except for the lack of tests) that's capable of doing 30,000-50,000 requests per second. Hopefully, you've gotten a feel for how Dropwizard combines Jetty, Jersey, Jackson, and other stable, mature libraries to provide a phenomenal platform for developing RESTful web applications.

There's a lot more to Dropwizard than is covered here (commands, bundles, servlets, advanced configuration, validation, HTTP clients, database clients, views, etc.), all of which is covered by the User Manual.

DROPWIZARD CORE

The `dropwizard-core` module provides you with everything you'll need for most of your applications.

It includes:

- Jetty, a high-performance HTTP server.
- Jersey, a full-featured RESTful web framework.
- Jackson, the best JSON library for the JVM.
- Metrics, an excellent library for application metrics.
- Logback, the successor to Log4j, Java's most widely-used logging framework.
- Hibernate Validator, the reference implementation of the Java Bean Validation standard.

Dropwizard consists mostly of glue code to automatically connect and configure these components.

4.1 Organizing Your Project

If you plan on developing a client library for other developers to access your service, we recommend you separate your projects into three Maven modules: `project-api`, `project-client`, and `project-application`.

`project-api` should contain your *Representations*; `project-client` should use those classes and an *HTTP client* to implement a full-fledged client for your application, and `project-application` should provide the actual application implementation, including *Resources*.

To give a concrete example of this project structure, let's say we wanted to create a *Stripe*-like API where clients can issue charges and the server would echo the charge back to the client. `stripe-api` project would hold our *Charge* object as both the server and client want to work with the charge and to promote code reuse, *Charge* objects are stored in a shared module. `stripe-app` is the Dropwizard application. `stripe-client` abstracts away the raw HTTP interactions and deserialization logic. Instead of using a HTTP client, users of `stripe-client` would just pass in a *Charge* object to a function and behind the scenes, `stripe-client` will call the HTTP endpoint. The client library may also take care of connection pooling, and may provide a more friendly way of interpreting error messages. Basically, distributing a client library for your app will help other developers integrate more quickly with the service.

If you are not planning on distributing a client library for developers, one can combine `project-api` and `project-application` into a single project, which tends to look like this:

- `com.example.myapplication`:
 - `api`: *Representations*. Request and response bodies.
 - `cli`: *Commands*
 - `client`: *Client* code that accesses external HTTP services.

- core: Domain implementation; where objects not used in the API such as POJOs, validations, crypto, etc, reside.
- db: *Database* access classes
- health: *Health Checks*
- resources: *Resources*
- MyApplication: The *application* class
- MyApplicationConfiguration: *configuration* class

4.2 Application

The main entry point into a Dropwizard application is, unsurprisingly, the `Application` class. Each `Application` has a **name**, which is mostly used to render the command-line interface. In the constructor of your `Application` you can add *Bundles* and *Commands* to your application.

4.3 Configuration

Dropwizard provides a number of built-in configuration parameters. They are well documented in the [example project's configuration](#) and [configuration reference](#).

Each `Application` subclass has a single type parameter: that of its matching `Configuration` subclass. These are usually at the root of your application's main package. For example, your `User` application would have two classes: `UserApplicationConfiguration`, extending `Configuration`, and `UserApplication`, extending `Application<UserApplicationConfiguration>`.

When your application runs *Configured Commands* like the `server` command, Dropwizard parses the provided YAML configuration file and builds an instance of your application's configuration class by mapping YAML field names to object field names.

Note: If your configuration file doesn't end in `.yaml` or `.yml`, Dropwizard tries to parse it as a JSON file.

To keep your configuration file and class manageable, we recommend grouping related configuration parameters into independent configuration classes. If your application requires a set of configuration parameters in order to connect to a message queue, for example, we recommend that you create a new `MessageQueueFactory` class:

Listing 1: Source file

```
public final class MessageQueueFactory {
    @NotEmpty
    private String host;

    @Min(1)
    @Max(65535)
    private int port = 5672;

    @JsonProperty
    public String getHost() {
        return host;
    }
}
```

(continues on next page)

(continued from previous page)

```
@JsonProperty
public void setHost(String host) {
    this.host = host;
}

@JsonProperty
public int getPort() {
    return port;
}

@JsonProperty
public void setPort(int port) {
    this.port = port;
}

public MessageQueueClient build(Environment environment) {
    MessageQueueClient client = new MessageQueueClient(getHost(), getPort());
    environment.lifecycle().manage(new Managed() {
        @Override
        public void start() {
        }

        @Override
        public void stop() {
            client.close();
        }
    });
    return client;
}
}
```

In this example our factory will automatically tie our `MessageQueueClient` connection to the lifecycle of our application's `Environment`.

Your main Configuration subclass can then include this as a member field:

Listing 2: Source file

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private MessageQueueFactory messageQueue = new MessageQueueFactory();

    @JsonProperty("messageQueue")
    public MessageQueueFactory getMessageQueueFactory() {
        return messageQueue;
    }

    @JsonProperty("messageQueue")
    public void setMessageQueueFactory(MessageQueueFactory factory) {
        this.messageQueue = factory;
    }
}
```

(continues on next page)

(continued from previous page)

```
}
```

And your `Application` subclass can then use your factory to directly construct a client for the message queue:

Listing 3: Source file

```
public void run(ExampleConfiguration configuration,
                Environment environment) {
    MessageQueueClient messageQueue = configuration.getMessageQueueFactory().
    ↪ build(environment);
}
```

Then, in your application's YAML file, you can use a nested `messageQueue` field:

Listing 4: Source file

```
messageQueue:
  host: mq.example.com
  port: 5673
```

The `@NotNull`, `@NotEmpty`, `@Min`, `@Max`, and `@Valid` annotations are part of *Dropwizard Validation* functionality. If your YAML configuration file's `messageQueue.host` field was missing (or was a blank string), Dropwizard would refuse to start and would output an error message describing the issues.

Once your application has parsed the YAML file and constructed its `Configuration` instance, Dropwizard then calls your `Application` subclass to initialize your application's `Environment`.

Note: You can override configuration settings by passing special Java system properties when starting your application. Overrides must start with prefix `dw.`, followed by the path to the configuration value being overridden.

For example, to override the Logging level, you could start your application like this:

```
java -Ddw.logging.level=DEBUG server my-config.json
```

This will work even if the configuration setting in question does not exist in your config file, in which case it will get added.

You can override configuration settings in arrays of objects like this:

```
java -Ddw.server.applicationConnectors[0].port=9090 server my-config.json
```

You can override configuration settings in maps like this:

```
java -Ddw.database.properties.hibernate.hbm2ddl.auto=none server my-config.json
```

If you need to use the `'` character in one of the values, you can escape it by using `\.` instead.

You can also override a configuration setting that is an array of strings by using the `'` character as an array element separator. For example, to override a configuration setting `myapp.myserver.hosts` that is an array of strings in the configuration, you could start your service like this: `java -Ddw.myapp.myserver.hosts=server1,server2,server3 server my-config.json`

If you need to use the `'` character in one of the values, you can escape it by using `\.` instead.

The array override facility only handles configuration elements that are arrays of simple strings. Also, the setting in question must already exist in your configuration file as an array; this mechanism will not work if the configuration key being overridden does not exist in your configuration file. If it does not exist or is not an array setting, it will get added as a simple string setting, including the `'` characters as part of the string.

4.3.1 Environment variables

The `dropwizard-configuration` module also provides the capabilities to substitute configuration settings with the value of environment variables using a `SubstitutingSourceProvider` and `EnvironmentVariableSubstitutor`.

Listing 5: Source file

```
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    // Enable variable substitution with environment variables
    EnvironmentVariableSubstitutor substitutor = new
↳EnvironmentVariableSubstitutor(false);
    SubstitutingSourceProvider provider =
        new SubstitutingSourceProvider(bootstrap.getConfiguratioSourceProvider(),
↳substitutor);
    bootstrap.setConfigurationSourceProvider(provider);
}
```

The configuration settings which should be substituted need to be explicitly written in the configuration file and follow the substitution rules of `StringSubstitutor` from the Apache Commons Text library.

Listing 6: Source file

```
mySetting: ${DW_MY_SETTING}
defaultSetting: ${DW_DEFAULT_SETTING:-default value}
```

In general `SubstitutingSourceProvider` isn't restricted to substitute environment variables but can be used to replace variables in the configuration source with arbitrary values by passing a custom `StringSubstitutor` implementation.

4.3.2 SSL

SSL support is built into Dropwizard. You will need to provide your own java keystore, which is outside the scope of this document (`keytool` is the command you need, and [Jetty's documentation](#) can get you started). There is a test keystore you can use in the [Dropwizard example project](#).

Listing 7: Source file

```
---
server:
  applicationConnectors:
    - type: https
      port: 8443
      keyStorePath: example.keystore
      keyStorePassword: example
      validateCerts: false
```

By default, only secure TLSv1.2 cipher suites are allowed. Older versions of cURL, Java 6 and 7, and other clients may be unable to communicate with the allowed cipher suites, but this was a conscious decision that sacrifices interoperability for security.

Dropwizard allows a workaround by specifying a customized list of cipher suites. If no lists of supported protocols or cipher suites are specified, then the JVM defaults are used. If no lists of excluded protocols or cipher suites are specified, then the defaults are inherited from Jetty.

The following list of excluded cipher suites will allow for TLSv1 and TLSv1.1 clients to negotiate a connection similar to pre-Dropwizard 1.0.

Listing 8: Source file

```
---
server:
  applicationConnectors:
    - type: https
      port: 8443
      excludedCipherSuites:
        - SSL_RSA_WITH_DES_CBC_SHA
        - SSL_DHE_RSA_WITH_DES_CBC_SHA
        - SSL_DHE_DSS_WITH_DES_CBC_SHA
        - SSL_RSA_EXPORT_WITH_RC4_40_MD5
        - SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
        - SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
        - SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
```

Since the version 9.4.8 (Dropwizard 1.2.3) Jetty supports native SSL via Google's [Conscrypt](#) which uses [BoringSSL](#) (Google's fork of OpenSSL) for handling cryptography. You can enable it in Dropwizard by registering the provider in your app:

Listing 9: Source file

```
<dependency>
  <groupId>org.conscrypt</groupId>
  <artifactId>conscrypt-openjdk-uber</artifactId>
  <version>${conscrypt.version}</version>
</dependency>
```

Listing 10: Source file

```
static {
  Security.insertProviderAt(new OpenSSLProvider(), 1);
}
```

and setting the JCE provider in the configuration:

Listing 11: Source file

```
---
server:
  type: simple
  connector:
    type: https
    jceProvider: Conscrypt
```

For HTTP/2 servers you need to add an ALPN Conscrypt provider as a dependency.

Listing 12: Source file

```
<dependency>
  <groupId>org.eclipse.jetty</groupId>
  <artifactId>jetty-alpn-conscrypt-server</artifactId>
```

(continues on next page)

(continued from previous page)

</dependency>

Note: If you are using Conscrypt with Java 8, you must exclude TLSv1.3 protocol as it is now enabled per default with Conscrypt 2.0.0 but not supported by Java 8.

4.4 Bootstrapping

Before a Dropwizard application can provide the command-line interface, parse a configuration file, or run as a server, it must first go through a bootstrapping phase. This phase corresponds to your `Application` subclass's `initialize` method. You can add *Bundles*, *Commands*, or register Jackson modules to allow you to include custom types as part of your configuration class.

4.5 Environments

A Dropwizard `Environment` consists of all the *Resources*, servlets, filters, *Health Checks*, *Health*, Jersey providers, *Managed Objects*, *Tasks*, and Jersey properties which your application provides.

Each `Application` subclass implements a `run` method. This is where you should be creating new resource instances, etc., and adding them to the given `Environment` class:

```
@Override
public void run(ExampleConfiguration config,
               Environment environment) {
    // encapsulate complicated setup logic in factories
    final Thingy thingy = config.getThingyFactory().build();

    environment.jersey().register(new ThingyResource(thingy));
    environment.healthChecks().register("thingy", new ThingyHealthCheck(thingy));
}
```

It's important to keep the `run` method clean, so if creating an instance of something is complicated, like the `Thingy` class above, extract that logic into a factory.

4.6 Health Checks

A health check is a runtime test which you can use to verify your application's behavior in its production environment. For example, you may want to ensure that your database client is connected to the database:

Listing 13: Source file

```
public class DatabaseHealthCheck extends HealthCheck {
    private final Database database;

    public DatabaseHealthCheck(Database database) {
        this.database = database;
    }
}
```

(continues on next page)

(continued from previous page)

```

@Override
protected Result check() throws Exception {
    if (database.isConnected()) {
        return Result.healthy();
    } else {
        return Result.unhealthy("Cannot connect to " + database.getUrl());
    }
}
}

```

You can then add this health check to your application's environment:

Listing 14: Source file

```
environment.healthChecks().register("database", new DatabaseHealthCheck(database));
```

By sending a GET request to /healthcheck on the admin port you can run these tests and view the results:

```

$ curl http://dw.example.com:8081/healthcheck
{"deadlocks":{"healthy":true},"database":{"healthy":true}}

```

If all health checks report success, a 200 OK is returned. If any fail, a 500 Internal Server Error is returned with the error messages and exception stack traces (if an exception was thrown).

Note: This behavior overlaps in many ways with the new *Health* functionality. If you wish to disable the admin health servlet, a new flag was introduced into the health check configuration *Health checks* to allow disabling it.

All Dropwizard applications ship with the *deadlocks* health check installed by default, which uses Java 1.6's built-in thread deadlock detection to determine if any threads are deadlocked.

4.7 Health

The health checks described in *Health Checks* can be configured to create a holistic view of your service health, which can then be used to drive decision making by things like *Kubernetes readiness & liveness checks*, or to dictate whether or not a load balancer should forward traffic to your service.

This can be done by running these dependency health checks periodically in the background on some schedule, and then aggregating the results of all of those checks into a single indicator of overall health. Certain dependencies may be critical to your application functioning, like a database that your service can't function without, but other dependencies may be more non-critical to your service being able to function (let's say a cache, that could be considered more of a nice to have than a necessity).

Define the following health check configurations in your *config.yml* file:

```

health:
  delayedShutdownHandlerEnabled: true
  shutdownWaitPeriod: 10s
  healthChecks:
    - name: user-database
      critical: true

```

(continues on next page)

(continued from previous page)

```

- name: user-notifications-queue
  critical: false
  schedule:
    checkInterval: 2500ms
    downtimeInterval: 10s
    failureAttempts: 2
    successAttempts: 1
- name: user-cache
  critical: false

```

Note: This behavior was integrated from the [Dropwizard Health](#) module. If you are migrating from that module to the new Dropwizard core framework health code, you will want to refer to *Migrating from dropwizard-health* the migration guide.

4.7.1 Application Status

There are two types of status that are supported currently: Alive and Ready.

- An **alive** status indicates the application is operating normally and does not need to be restarted to recover from a stuck state. Long-running applications can eventually reach a broken state and cannot recover except by being restarted (e.g. deadlocked threads).
- A **ready** status indicates the application is ready to serve traffic. Applications can temporarily be unable to serve traffic due to a variety of reasons, for example, an application might need to build/compute large caches during startup or can critically depend on an external service.

An example of how you might query the health check, assuming you're using the default responder/responseProvider settings in configuration:

```
https://<hostname>:<port>/health-check?type=<type>&name=<name>
```

- replace `<type>` with `ready` or `alive`; defaults to `ready` if the `type` parameter is not provided
- replace `<name>` with the name of the health check to query. Multiple names can be provided, or no no names. If all checks are desired, `name=all` can be specified to retrieve all checks

4.7.2 HTTP & TCP Checks

Should your service have any dependencies that it needs to perform health checks against that expose either an HTTP or TCP health check interface, you can use the `HttpHealthCheck` or `TcpHealthCheck` classes to do so easily.

You will need to register your health check(s) in your `Application` class `run()` method.

HTTP

```

@Override
public void run(final AppConfiguration configuration, final Environment environment) {
    ...
    environment.healthChecks().register("some-http-dependency", new HttpHealthCheck(
        ↪ "http://some-http-dependency.com:8080/health-check"));
}

```

TCP

```
@Override
public void run(final AppConfiguration configuration, final Environment environment) {
    ...
    environment.healthChecks().register("some-tcp-dependency", new TcpHealthCheck("some-
↪tcp-dependency.com", 443));
}
```

4.7.3 Health Data Access

In the *Application.run()* method, you can access views of the health state data in two different ways:

Accessing data directly

```
@Override
public void run(final AppConfiguration configuration, final Environment environment) {
    ...
    Collection<HealthStateView> views = environment.health().healthStateAggregator().
↪healthStateViews();
}
```

Listening to data changes

```
@Override
public void run(final AppConfiguration configuration, final Environment environment) {
    ...
    HealthStateListener myListener = new HealthStateListener() {
        @Override
        public void onStateChanged(String healthCheckName, boolean healthy) {
            System.out.println(healthCheckName + "changed state to " + healthy);
        }

        @Override
        public void onHealthyCheck(String healthCheckName) {
            System.out.println(healthCheckName + "is healthy! :)");
        }

        @Override
        public void onUnhealthyCheck(String healthCheckName) {
            System.out.println(healthCheckName + "is unhealthy! :)");
        }
    };
    environment.health().addHealthStateListener(myListener);
}
```

4.8 Managed Objects

Most applications involve objects which need to be started and stopped: thread pools, database connections, etc. Dropwizard provides the `Managed` interface for this. You can either have the class in question implement the `#start()` and/or `#stop()` methods, or write a wrapper class which does so. Adding a `Managed` instance to your application's `Environment` ties that object's lifecycle to that of the application's HTTP server. Before the server starts, the `#start()` method is called. After the server has stopped (and after its graceful shutdown period) the `#stop()` method is called.

For example, given a theoretical `Riak` client which needs to be started and stopped:

Listing 15: Source file

```
public class RiakClientManager implements Managed {
    private final RiakClient client;

    public RiakClientManager(RiakClient client) {
        this.client = client;
    }

    @Override
    public void start() throws Exception {
        client.start();
    }

    @Override
    public void stop() throws Exception {
        client.stop();
    }
}
```

Listing 16: Source file

```
public class ManagedApp extends Application<Configuration> {
    @Override
    public void run(Configuration configuration, Environment environment) {
        RiakClient client = new RiakClient();
        RiakClientManager riakClientManager = new RiakClientManager(client);
        environment.lifecycle().manage(riakClientManager);
    }
}
```

If `RiakClientManager#start()` throws an exception—e.g., an error connecting to the server—your application will not start and a full exception will be logged. If `RiakClientManager#stop()` throws an exception, the exception will be logged but your application will still be able to shut down.

It should be noted that `Environment` has built-in factory methods for `ExecutorService` and `ScheduledExecutorService` instances which are managed. These managed instances use `InstrumentedThreadFactory` that monitors the number of threads created, running and terminated

Listing 17: Source file

```
public void run(Configuration configuration, Environment environment) {
    ExecutorService executorService = environment.lifecycle()
        .executorService(nameFormat)
```

(continues on next page)

(continued from previous page)

```

        .maxThreads(maxThreads)
        .build();

    ScheduledExecutorService scheduledExecutorService = environment.lifecycle()
        .scheduledExecutorService(nameFormat)
        .build();
}

```

4.9 Bundles

A Dropwizard bundle is a reusable group of functionality, used to define blocks of an application's behavior by implementing the `ConfiguredBundle` interface.

For example, `AssetBundle` from the `dropwizard-assets` module provides a simple way to serve static assets from your application's `src/main/resources/assets` directory as files available from `/assets/*` (or any other path) in your application.

Given the bundle `MyConfiguredBundle` and the interface `MyConfiguredBundleConfig` below, your application's `Configuration` subclass would need to implement `MyConfiguredBundleConfig`.

Listing 18: Source file

```

public class MyConfiguredBundle implements ConfiguredBundle<MyConfiguredBundleConfig> {
    @Override
    public void run(MyConfiguredBundleConfig applicationConfig, Environment environment)
    ↪ {
        applicationConfig.getBundleSpecificConfig();
    }

    @Override
    public void initialize(Bootstrap<?> bootstrap) {
    }
}

```

Listing 19: Source file

```

public interface MyConfiguredBundleConfig {
    String getBundleSpecificConfig();
}

```

4.9.1 Serving Assets

Either your application or your static assets can be served from the root path, but not both. The latter is useful when using Dropwizard to back a Javascript application. To enable it, move your application to a sub-URL.

```

server:
  rootPath: /api/

```

Note: If you use the *Simple* server configuration, then `rootPath` is calculated relatively from

applicationContextPath. So, your API will be accessible from the path /application/api/

Then use an extended `AssetsBundle` constructor to serve resources in the `assets` folder from the root path. `index.htm` is served as the default page.

Listing 20: Source file

```
bootstrap.addBundle(new AssetsBundle("/assets/", "/"));
```

When an `AssetBundle` is added to the application, it is registered as a servlet using a default name of `assets`. If the application needs to have multiple `AssetBundle` instances, the extended constructor should be used to specify a unique name for the `AssetBundle`.

Listing 21: Source file

```
bootstrap.addBundle(new AssetsBundle("/assets/css", "/css", null, "css"));
bootstrap.addBundle(new AssetsBundle("/assets/js", "/js", null, "js"));
bootstrap.addBundle(new AssetsBundle("/assets/fonts", "/fonts", null, "fonts"));
```

4.9.2 SSL Reload

By registering the `SslReloadBundle` your application can have new certificate information reloaded at runtime, so a restart is not necessary.

Listing 22: Source file

```
public void initialize(Bootstrap<Configuration> bootstrap) {
    bootstrap.addBundle(new SslReloadBundle());
}
```

To trigger a reload send a POST request to `ssl-reload`

```
curl -k -X POST 'https://localhost:<admin-port>/tasks/ssl-reload'
```

Dropwizard will use the same exact https configuration (keystore location, password, etc) when performing the reload.

Note: If anything is wrong with the new certificate (eg. wrong password in keystore), no new certificates are loaded. So if the application and admin ports use different certificates and one of them is invalid, then none of them are reloaded.

A http 500 error is returned on reload failure, so make sure to trap for this error with whatever tool is used to trigger a certificate reload, and alert the appropriate admin. If the situation is not remedied, next time the app is stopped, it will be unable to start!

4.10 Commands

Commands are basic actions which Dropwizard runs based on the arguments provided on the command line. The built-in `server` command, for example, spins up an HTTP server and runs your application. Each `Command` subclass has a name and a set of command line options which Dropwizard will use to parse the given command line arguments.

Below is an example on how to add a command and have Dropwizard recognize it.

Listing 23: Source file

```
public class MyCommand extends Command {
    public MyCommand() {
        // The name of our command is "hello" and the description printed is
        // "Prints a greeting"
        super("hello", "Prints a greeting");
    }

    @Override
    public void configure(Subparser subparser) {
        // Add a command line option
        subparser.addArgument("-u", "--user")
            .dest("user")
            .type(String.class)
            .required(true)
            .help("The user of the program");
    }

    @Override
    public void run(Bootstrap<?> bootstrap, Namespace namespace) throws Exception {
        System.out.println("Hello " + namespace.getString("user"));
    }
}
```

Dropwizard recognizes our command once we add it in the `initialize` stage of our application.

Listing 24: Source file

```

public class CustomCommandApp extends Application<Configuration> {
    @Override
    public void initialize(Bootstrap<Configuration> bootstrap) {
        bootstrap.addCommand(new MyCommand());
    }

    @Override
    public void run(Configuration configuration, Environment environment) {
    }
}

```

To invoke the new functionality, run the following:

```
java -jar <jarfile> hello dropwizard
```

4.10.1 Configured Commands

Some commands require access to configuration parameters and should extend the `ConfiguredCommand` class, using your application's `Configuration` class as its type parameter. By default, Dropwizard will treat the last argument on the command line as the path to a YAML configuration file, parse and validate it, and provide your command with an instance of the configuration class.

A `ConfiguredCommand` can have additional command line options specified, while keeping the last argument the path to the YAML configuration.

Listing 25: Source file

```

public class MyConfiguredCommand extends ConfiguredCommand<Configuration> {
    public MyConfiguredCommand() {
        // The name of our command is "hello" and the description printed is
        // "Prints a greeting"
        super("hello", "Prints a greeting");
    }

    @Override
    public void configure(Subparser subparser) {
        // Add a command line option
        subparser.addArgument("-u", "--user")
            .dest("user")
            .type(String.class)
            .required(true)
            .help("The user of the program");
    }

    @Override
    protected void run(Bootstrap<Configuration> bootstrap,
        Namespace namespace,
        Configuration configuration) throws Exception {
        System.out.println("Hello " + namespace.getString("user"));
    }
}

```

For more advanced customization of the command line (for example, having the configuration file location specified by `-c`), adapt the `ConfiguredCommand` class as needed.

Note: If you override the `configure` method, you **must** call `super.override(subparser)` (or call `addFileArgument`) in order to preserve the configuration file parameter in the subparser.

4.11 Tasks

A Task is a run-time action your application provides access to on the administrative port via HTTP. All Dropwizard applications start with: the `gc` task, which explicitly triggers the JVM's garbage collection (This is useful, for example, for running full garbage collections during off-peak times or while the given application is out of rotation.); and the `log-level` task, which configures the level of any number of loggers at runtime (akin to Logback's `JmxConfigurator`). The `execute` method of a Task can be annotated with `@Timed`, `@Metered`, and `@ExceptionMetered`. Dropwizard will automatically record runtime information about your tasks. Here's a basic task class:

Listing 26: Source file

```
public class TruncateDatabaseTask extends Task {
    private final Database database;

    public TruncateDatabaseTask(Database database) {
        super("truncate");
        this.database = database;
    }

    @Override
    public void execute(Map<String, List<String>> parameters, PrintWriter output) throws
↳Exception {
        this.database.truncate();
    }
}
```

You can then add this task to your application's environment:

Listing 27: Source file

```
public void run(Configuration configuration, Environment environment) {
    Database database = new Database();
    environment.admin().addTask(new TruncateDatabaseTask(database));
}
```

Running a task can be done by sending a POST request to `/tasks/{task-name}` on the admin port. The task will receive any query parameters as arguments. For example:

```
$ curl -X POST http://dw.example.com:8081/tasks/gc
Running GC...
Done!
```

You can also extend `PostBodyTask` to create a task which uses the body of the post request. Here's an example:

Listing 28: Source file

```

public class EchoTask extends PostBodyTask {
    public EchoTask() {
        super("echo");
    }

    @Override
    public void execute(Map<String, List<String>> parameters, String body, PrintWriter
↵output) throws Exception {
        output.write(body);
        output.flush();
    }
}

```

4.12 Logging

Dropwizard uses [Logback](#) for its logging backend. It provides an [slf4j](#) implementation, and even routes all `java.util.logging`, `Log4j`, and Apache Commons Logging usage through Logback.

slf4j provides the following logging levels:

ERROR

Error events that might still allow the application to continue running.

WARN

Potentially harmful situations.

INFO

Informational messages that highlight the progress of the application at coarse-grained level.

DEBUG

Fine-grained informational events that are most useful to debug an application.

TRACE

Finer-grained informational events than the DEBUG level.

Note: If you don't want to use Logback, you can exclude it from Dropwizard and use an alternative logging configuration:

- Exclude Logback from the `dropwizard-core` artifact

```

<dependency>
  <groupId>io.dropwizard</groupId>
  <artifactId>dropwizard-core</artifactId>
  <version>${dropwizard.version}</version>
  <exclusions>
    <exclusion>
      <groupId>ch.qos.logback</groupId>
      <artifactId>logback-classic</artifactId>
    </exclusion>
    <exclusion>
      <groupId>ch.qos.logback</groupId>
      <artifactId>logback-access</artifactId>
    </exclusion>
  </exclusions>
</dependency>

```

(continues on next page)

(continued from previous page)

```
</exclusion>
<exclusion>
  <groupId>org.slf4j</groupId>
  <artifactId>log4j-over-slf4j</artifactId>
</exclusion>
</exclusions>
</dependency>
```

- Mark the logging configuration as external in your Dropwizard config

```
server:
  type: simple
  applicationContextPath: /application
  adminContextPath: /admin
  requestLog:
    type: external
  logging:
    type: external
```

- Disable bootstrapping Logback in your application

```
public class ExampleApplication extends Application<ExampleConfiguration> {

    @Override
    protected void bootstrapLogging() {
    }
}
```

4.12.1 Log Format

Dropwizard's log format has a few specific goals:

- Be human readable.
- Be machine parsable.
- Be easy for sleepy ops folks to figure out why things are pear-shaped at 3:30AM using standard UNIXy tools like tail and grep.

The logging output looks like this:

```
TRACE [2010-04-06 06:42:35,271] com.example.dw.Thing: Contemplating doing a thing.
DEBUG [2010-04-06 06:42:35,274] com.example.dw.Thing: About to do a thing.
INFO [2010-04-06 06:42:35,274] com.example.dw.Thing: Doing a thing
WARN [2010-04-06 06:42:35,275] com.example.dw.Thing: Doing a thing
ERROR [2010-04-06 06:42:35,275] com.example.dw.Thing: This may get ugly.
! java.lang.RuntimeException: oh noes!
! at com.example.dw.Thing.run(Thing.java:16)
!
```

A few items of note:

- All timestamps are in UTC and ISO 8601 format.

- You can grep for messages of a specific level really easily:

```
tail -f dw.log | grep '^WARN'
```

- You can grep for messages from a specific class or package really easily:

```
tail -f dw.log | grep 'com.example.dw.Thing'
```

- You can even pull out full exception stack traces, plus the accompanying log message:

```
tail -f dw.log | grep -B 1 '^\!'
```

- The `!` prefix does *not* apply to syslog appenders, as stack traces are sent separately from the main message. Instead, `t` is used (this is the default value of the `SyslogAppender` that comes with Logback). This can be configured with the `stackTracePrefix` option when defining your appender.

To apply the prefixing with the `!` symbol Dropwizard introduces several new [Logback conversion words](#). These are `dwEx`, `dwException`, `dwThrowable`, `dwXEx`, `dwXException`, `dwXThrowable`, `dwREx` and `dwRootException`. These conversion words work like the ones from Logback, except that the first tab of a stack trace is replaced by a `!`.

The default Dropwizard logging layout uses the Dropwizard specific conversion words.

4.12.2 Configuration

You can specify a default logger level, override the levels of other loggers in your YAML configuration file, and even specify appenders for them. The latter form of configuration is preferable, but the former is also acceptable.

```
# Logging settings.
logging:

  # The default level of all loggers. Can be OFF, ERROR, WARN, INFO, DEBUG, TRACE, or
  ↪ ALL.
  level: INFO

  # Logger-specific levels.
  loggers:

    # Overrides the level of com.example.dw.Thing and sets it to DEBUG.
    "com.example.dw.Thing": DEBUG

    # Enables the SQL query log and redirect it to a separate file
    "org.hibernate.SQL":
      level: DEBUG
      # This line stops org.hibernate.SQL (or anything under it) from using the root
      ↪ logger
      additive: false
      appenders:
        - type: file
          currentLogFilename: ./logs/example-sql.log
          archivedLogFilenamePattern: ./logs/example-sql-%d.log.gz
          archivedFileCount: 5
```

4.12.3 Asynchronous Logging

By default, all logging in Dropwizard is asynchronous, even to typically synchronous sinks such as files and the console. When a slow logger (like file logger on an overloaded disk) is coupled with a high load, Dropwizard will seamlessly drop events of lower importance (TRACE, DEBUG, INFO) in an attempt to maintain reasonable latency.

Tip: Instead of logging business critical statements under INFO, insert them into a database, durable message queue, or another mechanism that gives confidence that the request has satisfied business requirements before returning the response to the client.

This logging behavior *can be configured*:

- Set `discardingThreshold` to 0 so that no events are dropped
- At the opposite end, set `neverBlock` to `true` so that even WARN and ERROR levels will be discarded from logging under heavy load

Request access logging has the same logging behavior, and since all request logging is done under INFO, each log statement has an equal chance of being dropped if the log queue is nearing full.

4.12.4 Console Logging

By default, Dropwizard applications log INFO and higher to STDOUT. You can configure this by editing the logging section of your YAML configuration file:

```
logging:
  appenders:
    - type: console
      threshold: WARN
      target: stderr
```

In the above, we're instead logging only WARN and ERROR messages to the STDERR device.

4.12.5 File Logging

Dropwizard can also log to an automatically rotated set of log files. This is the recommended configuration for your production environment:

```
logging:
  appenders:
    - type: file
      # The file to which current statements will be logged.
      currentLogFilename: ./logs/example.log

      # When the log file rotates, the archived log will be renamed to this and gzipped.
      ↪The
      # %d is replaced with the previous day (yyyy-MM-dd). Custom rolling windows can be
      ↪created
      # by passing a SimpleDateFormat-compatible format as an argument: "%d{yyyy-MM-dd-
      ↪hh}".
      archivedLogFilenamePattern: ./logs/example-%d.log.gz
```

(continues on next page)

(continued from previous page)

```

# The number of archived files to keep.
archivedFileCount: 5

# The timezone used to format dates. HINT: USE THE DEFAULT, UTC.
timeZone: UTC

```

4.12.6 Syslog Logging

Finally, Dropwizard can also log statements to syslog.

Note: Because Java doesn't use the native syslog bindings, your syslog server **must** have an open network socket.

logging:

```

appenders:
- type: syslog
  # The hostname of the syslog server to which statements will be sent.
  # N.B.: If this is the local host, the local syslog instance will need to be
  ↪configured to
  # listen on an inet socket, not just a Unix socket.
  host: localhost

  # The syslog facility to which statements will be sent.
  facility: local0

```

You can combine any number of different **appenders**, including multiple instances of the same appender with different configurations:

logging:

```

# Permit DEBUG, INFO, WARN and ERROR messages to be logged by appenders.
level: DEBUG

appenders:
  # Log warnings and errors to stderr
  - type: console
    threshold: WARN
    target: stderr

  # Log info, warnings and errors to our apps' main log.
  # Rolled over daily and retained for 5 days.
  - type: file
    threshold: INFO
    currentLogFilename: ./logs/example.log
    archivedLogFilenamePattern: ./logs/example-%d.log.gz
    archivedFileCount: 5

  # Log debug messages, info, warnings and errors to our apps' debug log.

```

(continues on next page)

(continued from previous page)

```
# Rolled over hourly and retained for 6 hours
- type: file
  threshold: DEBUG
  currentLogFilename: ./logs/debug.log
  archivedLogFilenamePattern: ./logs/debug-%d{yyyy-MM-dd-hh}.log.gz
  archivedFileCount: 6
```

4.12.7 JSON Log Format

You may prefer to produce logs in a structured format such as JSON, so it can be processed by analytics or BI software. For that, add a module to the project for supporting JSON layouts:

```
<dependency>
  <groupId>io.dropwizard</groupId>
  <artifactId>dropwizard-json-logging</artifactId>
  <version>${dropwizard.version}</version>
</dependency>
```

Setup the JSON layout in the configuration file.

For general logging:

```
logging:
  appenders:
    - type: console
      layout:
        type: json
```

The json layout will produces the following log message:

```
{"timestamp":1515002688000, "level":"INFO", "logger":"org.eclipse.jetty.server.Server",
  → "thread":"main", "message":"Started @6505ms"}
```

For request logging:

```
server:
  requestLog:
    appenders:
      - type: console
        layout:
          type: access-json
```

The access-json layout will produces the following log message:

```
{"timestamp":1515002688000, "method":"GET", "uri":"/hello-world", "status":200, "protocol
  → ":"HTTP/1.1", "contentLength":37, "remoteAddress":"127.0.0.1", "requestTime":5, "userAgent
  → ":"Mozilla/5.0"}
```

4.12.8 Logging Configuration via HTTP

Active log levels can be changed during the runtime of a Dropwizard application via HTTP using the `LogConfigurationTask`. For instance, to configure the log level for a single `Logger`. The `logger` parameter may be repeated. The optional `duration` parameter must be an ISO 8601 duration format. When the duration elapses the level will revert to the effective level of the parent logger.:

```
# Configure com.example.helloworld to INFO
curl -X POST -d "logger=com.example.helloworld&level=INFO" http://localhost:8081/tasks/
↳log-level
# Configure com.example.helloworld and com.example.helloearth to INFO
curl -X POST -d "logger=com.example.helloworld&logger=com.example.helloearth&level=INFO"
↳http://localhost:8081/tasks/log-level
# Configure com.example.helloworld to INFO, then revert to default level after 10 minutes
curl -X POST -d "logger=com.example.helloworld&level=INFO&duration=PT10M" http://
↳localhost:8081/tasks/log-level
# Revert com.example.helloworld to the default level
curl -X POST -d "logger=com.example.helloworld" http://localhost:8081/tasks/log-level
```

Note: Chaining log level changes on the same package may have unexpected consequences due to the naive implementation of a simple FIFO timer.

4.12.9 Logging Filters

Just because a statement has a level of `INFO`, doesn't mean it should be logged with other `INFO` statements. One can create logging filters that will intercept log statements before they are written and decide if they're allowed. Log filters can work on both regular statements and request log statements. The following example will be for request logging as there are many reasons why certain requests may be excluded from the log:

- Only log requests that have large bodies
- Only log requests that are slow
- Only log requests that resulted in a non-2xx status code
- Exclude requests that contain sensitive information in the URL
- Exclude healthcheck requests

The example will demonstrate excluding `/secret` requests from the log.

```
@JsonTypeName("secret-filter-factory")
public class SecretFilterFactory implements FilterFactory<IAccessEvent> {
    @Override
    public Filter<IAccessEvent> build() {
        return new Filter<IAccessEvent>() {
            @Override
            public FilterReply decide(IAccessEvent event) {
                if (event.getRequestURI().equals("/secret")) {
                    return FilterReply.DENY;
                } else {
                    return FilterReply.NEUTRAL;
                }
            }
        }
    }
}
```

(continues on next page)

(continued from previous page)

```

    };
  }
}

```

Reference `SecretFilterFactory` type in our configuration.

```

server:
  requestLog:
    appenders:
      - type: console
      filterFactories:
        - type: secret-filter-factory

```

The last step is to add our class (in this case `com.example.SecretFilterFactory`) to `META-INF/services/io.dropwizard.logging.common.filter.FilterFactory` in our resources folder.

4.12.10 Filtering Request Logs for a Specific URI

Reference `UriFilterFactory` type in your configuration.

```

server:
  requestLog:
    appenders:
      - type: console
      filterFactories:
        - type: uri
        uris:
          - "/health-check"

```

4.13 Testing Applications

All of Dropwizard's APIs are designed with testability in mind, so even your applications can have unit tests:

Listing 29: Source file

```

@ExtendWith(MockitoExtension.class)
class MyApplicationTest {
    @Mock
    private Environment environment;
    @Mock
    private JerseyEnvironment jersey;
    private MyApplication application;
    private MyConfiguration config;

    @BeforeEach
    void setup() throws Exception {
        config = new MyConfiguration();
        config.setMyParam("yay");
        application = new MyApplication();
        when(environment.jersey()).thenReturn(jersey);
    }
}

```

(continues on next page)

(continued from previous page)

```

}

@Test
void buildsMyResource() throws Exception {
    application.run(config, environment);

    verify(jersey).register(eq(MyResource.class));
}
}

```

We highly recommend [Mockito](#) for all your mocking needs.

4.14 Banners

We think applications should print out a big ASCII art banner on startup. Yours should, too. It's fun. Just add a `banner.txt` class to `src/main/resources` and it'll print it out when your application starts:

```

INFO [2011-12-09 21:56:37,209] io.dropwizard.core.cli.ServerCommand: Starting hello-
↪world

                                dP
                                88
.d8888b. dP. .dP .d8888b. 88d8b.d8b. 88d888b. 88 .d8888b.
88oooo8 `8bd8' 88' `88 88'`88'`88 88' `88 88 88oooo8
88. ... .d88b. 88. .88 88 88 88. .88 88 88. ...
`88888P' dP' `dP `88888P8 dP dP dP 88Y888P' dP `88888P'
                                88
                                dP

INFO [2011-12-09 21:56:37,214] org.eclipse.jetty.server.Server: jetty-7.6.0
...

```

We could probably make up an argument about why this is a serious devops best practice with high ROI and an Agile Tool, but honestly we just enjoy this.

We recommend you use [TAAG](#) for all your ASCII art banner needs.

4.15 Resources

Unsurprisingly, most of your day-to-day work with a Dropwizard application will be in the resource classes, which model the resources exposed in your RESTful API. Dropwizard uses [Jersey](#) for this, so most of this section is just re-hashing or collecting various bits of Jersey documentation.

Jersey is a framework for mapping various aspects of incoming HTTP requests to POJOs and then mapping various aspects of POJOs to outgoing HTTP responses. Here's a basic resource class:

```

@Path("/{user}/notifications")
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
public class NotificationsResource {
    private final NotificationStore store;
}

```

(continues on next page)

(continued from previous page)

```

public NotificationsResource(NotificationStore store) {
    this.store = store;
}

@GET
public NotificationList fetch(@PathParam("user") OptionalLong userId,
                             @QueryParam("count") @DefaultValue("20") OptionalInt_
↪count) {
    final List<Notification> notifications = store.fetch(userId.get(), count.get());
    if (notifications != null) {
        return new NotificationList(userId, notifications);
    }
    throw new WebApplicationException(Status.NOT_FOUND);
}

@POST
public Response add(@PathParam("user") OptionalLong userId,
                   @NotNull @Valid Notification notification) {
    final long id = store.add(userId.get(), notification);
    return Response.created(UriBuilder.fromResource(NotificationResource.class)
                               .build(userId.get(), id))
        .build();
}
}

```

This class provides a resource (a user's list of notifications) which responds to GET and POST requests to `/user/notifications`, providing and consuming `application/json` representations. There's quite a lot of functionality on display here, and this section will explain in detail what's in play and how to use these features in your application.

4.15.1 Paths

Important: Every resource class must have a `@Path` annotation.

The `@Path` annotation isn't just a static string, it's a [URI Template](#). The `{user}` part denotes a named variable, and when the template matches a URI the value of that variable will be accessible via `@PathParam`-annotated method parameters.

For example, an incoming request for `/1001/notifications` would match the URI template, and the value `"1001"` would be available as the path parameter named `user`.

If your application doesn't have a resource class whose `@Path` URI template matches the URI of an incoming request, Jersey will automatically return a `404 Not Found` to the client.

4.15.2 Methods

Methods on a resource class which accept incoming requests are annotated with the HTTP methods they handle: @GET, @POST, @PUT, @DELETE, @HEAD, @OPTIONS, @PATCH.

Support for arbitrary new methods can be added via the @HttpMethod annotation. They also must be added to the *list of allowed methods*. This means, by default, methods such as CONNECT and TRACE are blocked, and will return a 405 Method Not Allowed response.

If a request comes in which matches a resource class's path but has a method which the class doesn't support, Jersey will automatically return a 405 Method Not Allowed to the client.

The return value of the method (in this case, a NotificationList instance) is then mapped to the *negotiated media type*. In this case, our resource only supports JSON, and so the NotificationList is serialized to JSON using Jackson.

4.15.3 Metrics

Every resource method or the class itself can be annotated with @Timed, @Metered, @ResponseMetered and @ExceptionMetered. If the annotation is placed on the class, it will apply to all its resource methods. Dropwizard augments Jersey to automatically record runtime information about your resource methods.

```
public class ExampleApplication extends ResourceConfig {
    .
    .
    .
    register(new InstrumentedResourceMethodApplicationListener (new MetricRegistry()));
    config = config.register(ExampleResource.class);
    .
    .
    .
}

@Path("/example")
@Produces(MediaType.TEXT_PLAIN)
public class ExampleResource {
    @GET
    @Timed
    public String show() {
        return "yay";
    }

    @GET
    @Metered(name = "fancyName") // If name isn't specified, the meter will given the
    ↪name of the method it decorates.
    @Path("/metered")
    public String metered() {
        return "woo";
    }

    @GET
    @ExceptionMetered(cause = IOException.class) // Default cause is Exception.class
    @Path("/exception-metered")
    public String exceptionMetered(@QueryParam("splode") @DefaultValue("false") boolean
```

(continues on next page)

(continued from previous page)

```

↪splode) throws IOException {
    if (splode) {
        throw new IOException("AUGH");
    }
    return "fuh";
}

@GET
@ResponseMetered
@Path("/response-metered")
public Response responseMetered(@QueryParam("invalid") @DefaultValue("false")
↪boolean invalid) {
    if (invalid) {
        return Response.status(Response.Status.INTERNAL_SERVER_ERROR).build();
    }
    return Response.ok().build();
}
}

```

- @Timed measures the duration of requests to a resource
- @Metered measures the rate at which the resource is accessed
- @ResponseMetered measures rate for each class of response codes (1xx/2xx/3xx/4xx/5xx)
- @ExceptionMetered measures how often exceptions occur processing the resource

Important: @Timed and @Metered can only be used on the same resource method at the same time, if their name is unique, also see the annotation parameter name. Otherwise, the generated metrics names will be identical which will cause an `IllegalArgumentException`.

4.15.4 Parameters

The annotated methods on a resource class can accept parameters which are mapped to from aspects of the incoming request.

For example:

- A `@PathParam("user")`-annotated `String` takes the raw value from the `user` variable in the matched URI template and passes it into the method as a `String`.
- A `@QueryParam("count")`-annotated `OptionalInt` parameter takes the first `count` value from the request's query string and passes it as a `String` to `OptionalInt`'s constructor. `OptionalInt` parses the string as an `Integer`, returning a `400 Bad Request` if the value is malformed.
- A `@FormParam("name")`-annotated `Set<String>` parameter takes all the `name` values from a posted form and passes them to the method as a set of strings.
- A `*Param`-annotated `NonEmptyStringParam` will interpret empty strings as absent strings, which is useful in cases where the endpoint treats empty strings and absent strings as interchangeable.

What's noteworthy here is that you can actually encapsulate the vast majority of your validation logic using specialized parameter objects. See `AbstractParam` for details.

4.15.5 Request Entities

If you're handling request entities (e.g., an `application/json` object on a PUT request), you can model this as a parameter without a `*Param` annotation. In the *example code*, the `add` method provides a good example of this:

```
@POST
public Response add(@PathParam("user") OptionalLong userId,
                   @NotNull @Valid Notification notification) {
    final long id = store.add(userId.get(), notification);
    return Response.created(UriBuilder.fromResource(NotificationResource.class)
                           .build(userId.get(), id))
                   .build();
}
```

Jersey maps the request entity to any single, unbound parameter. In this case, because the resource is annotated with `@Consumes(MediaType.APPLICATION_JSON)`, it uses the Dropwizard-provided Jackson support which, in addition to parsing the JSON and mapping it to an instance of `Notification`, also runs that instance through Dropwizard's *Constraining Entities*.

If the deserialized `Notification` isn't valid, Dropwizard returns a `422 Unprocessable Entity` response to the client.

Note: If a request entity parameter is just annotated with `@Valid`, it is still allowed to be `null`, so to ensure that the object is present and validated `@NotNull @Valid` is a powerful combination.

4.15.6 Media Types

Jersey also provides full content negotiation, so if your resource class consumes `application/json` but the client sends a `text/plain` entity, Jersey will automatically reply with a `406 Not Acceptable`. Jersey's even smart enough to use client-provided `q`-values in their `Accept` headers to pick the best response content type based on what both the client and server will support.

4.15.7 Responses

If your clients are expecting custom headers or additional information (or, if you simply desire an additional degree of control over your responses), you can return explicitly-built `Response` objects:

```
return Response.noContent().language(Locale.GERMAN).build();
```

In general, though, we recommend you return actual domain objects if at all possible. It makes *testing resources* much easier.

4.15.8 Error Handling

Almost as important as an application's happy path (receiving expected input and returning expected output) is an application's behavior when something goes wrong.

If your resource class unintentionally throws an exception, Dropwizard will log that exception under the `ERROR` level (including stack traces) and return a terse, safe `application/json 500 Internal Server Error` response. The response will contain an ID that can be grepped out the server logs for additional information.

If your resource class needs to return an error to the client (e.g., the requested record doesn't exist), you have two options: throw a subclass of `Exception` or restructure your method to return a `Response`. If at all possible, prefer throwing `Exception` instances to returning `Response` objects, as that will make resource endpoints more self describing and easier to test.

The least intrusive way to map error conditions to a response is to throw a `WebApplicationException`:

```
@GET
@Path("/{collection}")
public Saying reduceCols(@PathParam("collection") String collection) {
    if (!collectionMap.containsKey(collection)) {
        final String msg = String.format("Collection %s does not exist", collection);
        throw new WebApplicationException(msg, Status.NOT_FOUND)
    }

    // ...
}
```

In this example a GET request to `/foobar` will return

```
{"code":404,"message":"Collection foobar does not exist"}
```

One can also take exceptions that your resource may throw and map them to appropriate responses. For instance, an endpoint may throw `IllegalArgumentException` and it may be worthy enough of a response to warrant a custom metric to track how often the event occurs. Here's an example of such an `ExceptionHandler`

```
public class IllegalArgumentExceptionMapper implements ExceptionMapper
↳<IllegalArgumentException> {
    private final Meter exceptions;
    public IllegalArgumentExceptionMapper(MetricRegistry metrics) {
        exceptions = metrics.meter(name(getClass(), "exceptions"));
    }

    @Override
    public Response toResponse(IllegalArgumentException e) {
        exceptions.mark();
        return Response.status(Status.BAD_REQUEST)
            .header("X-YOU-SILLY", "true")
            .type(MediaType.APPLICATION_JSON_TYPE)
            .entity(new ErrorMessage(Status.BAD_REQUEST.getStatusCode(),
                "You passed an illegal argument!"))
            .build();
    }
}
```

and then registering the exception mapper:

```
@Override
public void run(final MyConfiguration conf, final Environment env) {
    env.jersey().register(new IllegalArgumentExceptionMapper(env.metrics()));
    env.jersey().register(new Resource());
}
```

Overriding Default Exception Mappers

To override a specific exception mapper, register your own class that implements the same `ExceptionHandler<T>` as one of the default. For instance, we can customize responses caused by Jackson exceptions:

```
public class JsonProcessingExceptionHandler implements ExceptionMapper
<JsonProcessingException> {
    @Override
    public Response toResponse(JsonProcessingException exception) {
        // create the response
    }
}
```

With this method, one doesn't need to know what the default exception mappers are, as they are overridden if the user supplies a conflicting mapper. While not preferential, one can also disable all default exception mappers, by setting `server.registerDefaultExceptionMappers` to `false`. See the class `ExceptionHandlerBinder` for a list of the default exception mappers.

4.15.9 URIs

While Jersey doesn't quite have first-class support for hyperlink-driven applications, the provided `UriBuilder` functionality does quite well.

Rather than duplicate resource URIs, it's possible (and recommended!) to initialize a `UriBuilder` with the path from the resource class itself:

```
UriBuilder.fromResource(UserResource.class).build(user.getId());
```

4.15.10 Testing

As with just about everything in Dropwizard, we recommend you design your resources to be testable. Dependencies which aren't request-injected should be passed in via the constructor and assigned to `final` fields.

Testing, then, consists of creating an instance of your resource class and passing it a mock. (Again: [Mockito](#).)

```
public class NotificationsResourceTest {
    private final NotificationStore store = mock(NotificationStore.class);
    private final NotificationsResource resource = new NotificationsResource(store);

    @Test
    public void getsReturnNotifications() {
        final List<Notification> notifications = mock(List.class);
        when(store.fetch(1, 20)).thenReturn(notifications);

        final NotificationList list = resource.fetch(new LongParam("1"), new IntParam("20
```

(continues on next page)

(continued from previous page)

```

↪"));

    assertThat(list.getUserId(),
               is(1L));

    assertThat(list.getNotifications(),
               is(notifications));
}
}

```

4.15.11 Caching

Adding a Cache-Control statement to your resource class is simple with Dropwizard:

```

@GET
@CacheControl(maxAge = 6, maxAgeUnit = TimeUnit.HOURS)
public String getCachableValue() {
    return "yay";
}

```

The @CacheControl annotation will take all of the parameters of the Cache-Control header.

4.15.12 Sessions

Although Dropwizard's main purpose is to build stateless RESTful APIs, a stateful web service can be built using HTTP sessions. As most users won't profit from having session support enabled by default, session support is implemented as opt-in.

The underlying Jetty server will handle sessions only if a SessionHandler is provided at application startup. Therefore the following code has to be added to the run method of the Application class:

```

@Override
public void run(final TestConfiguration configuration, final Environment environment) {
    environment.servlets().setSessionHandler(new org.eclipse.jetty.server.session.
↪SessionHandler());
}

```

This will provide Jetty's default SessionHandler to the servlet environment and session support is enabled. To get an HttpSession object injected into a Jersey resource method, Dropwizard provides a @Session annotation:

```

public Response doSomethingWithSessions(@Session HttpSession httpSession) {
    return Response.ok().build();
}

```


4.16 Representations

Representation classes are classes which, when handled to various Jersey `MessageBodyReader` and `MessageBodyWriter` providers, become the entities in your application's API. Dropwizard heavily favors JSON, but it's possible to map from any POJO to custom formats and back.

4.16.1 Basic JSON

Jackson is awesome at converting regular POJOs to JSON and back. This file:

```
public class Notification {
    private String text;

    public Notification(String text) {
        this.text = text;
    }

    @JsonProperty
    public String getText() {
        return text;
    }

    @JsonProperty
    public void setText(String text) {
        this.text = text;
    }
}
```

gets converted into this JSON:

```
{
  "text": "hey it's the value of the text field"
}
```

If, at some point, you need to change the JSON field name or the Java field without affecting the other, you can add an explicit field name to the `@JsonProperty` annotation.

If you prefer immutable objects rather than JavaBeans, that's also doable:

```
public class Notification {
    private final String text;

    @JsonCreator
    public Notification(@JsonProperty("text") String text) {
        this.text = text;
    }

    @JsonProperty("text")
    public String getText() {
        return text;
    }
}
```

4.16.2 Advanced JSON

Not all JSON representations map nicely to the objects your application deals with, so it's sometimes necessary to use custom serializers and deserializers. Just annotate your object like this:

```
@JsonSerialize(using=FunkySerializer.class)
@JsonDeserialize(using=FunkyDeserializer.class)
public class Funky {
    // ...
}
```

Then make a `FunkySerializer` class which implements `JsonSerializer<Funky>` and a `FunkyDeserializer` class which implements `JsonDeserializer<Funky>`.

Snake Case

A common issue with JSON is the disagreement between `camelCase` and `snake_case` field names. Java and Javascript folks tend to like `camelCase`; Ruby, Python, and Perl folks insist on `snake_case`. To make Dropwizard automatically convert field names to `snake_case` (and back), just annotate the class with `@JsonSnakeCase`:

```
@JsonSnakeCase
public class Person {
    private final String firstName;

    @JsonCreator
    public Person(@JsonProperty String firstName) {
        this.firstName = firstName;
    }

    @JsonProperty
    public String getFirstName() {
        return firstName;
    }
}
```

This gets converted into this JSON:

```
{
  "first_name": "Coda"
}
```

Unknown properties

If the name of a JSON property cannot be mapped to a Java property (or otherwise handled), that JSON property will simply be ignored.

You can change this behavior by configuring Dropwizard's object mapper:

```
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.getObjectMapper().enable(DeserializationFeature.FAIL_ON_UNKNOWN_
↪PROPERTIES);
}
```

Note: The YAML configuration parser will fail on unknown properties regardless of the object mapper configuration.

4.16.3 Streaming Output

If your application happens to return lots of information, you may get a big performance and efficiency bump by using streaming output. By returning an object which implements Jersey's `StreamingOutput` interface, your method can stream the response entity in a chunk-encoded output stream. Otherwise, you'll need to fully construct your return value and *then* hand it off to be sent to the client.

4.16.4 HTML Representations

For generating HTML pages, check out Dropwizard's *views support*.

4.16.5 Custom Representations

Sometimes, though, you've got some wacky output format you need to produce or consume and no amount of arguing will make JSON acceptable. That's unfortunate but OK. You can add support for arbitrary input and output formats by creating classes which implement Jersey's `MessageBodyReader<T>` and `MessageBodyWriter<T>` interfaces. (Make sure they're annotated with `@Provider` and `@Produces("text/gibberish")` or `@Consumes("text/gibberish")`.) Once you're done, just add instances of them (or their classes if they depend on Jersey's `@Context` injection) to your application's `Environment` on initialization.

4.17 Filters

There might be cases when you want to filter out requests or modify them before they reach your Resources.

4.17.1 Jersey filters

Jersey has a rich api for *filters* and *interceptors* that can be used directly in Dropwizard. You can stop the request from reaching your resources by throwing a `WebApplicationException`. Alternatively, you can use filters to modify inbound requests or outbound responses.

```
@Provider
public class DateNotSpecifiedFilter implements ContainerRequestFilter {
    @Override
    public void filter(ContainerRequestContext requestContext) throws IOException {
        String dateHeader = requestContext.getHeaderString(Headers.DATE);

        if (dateHeader == null) {
            Exception cause = new IllegalArgumentException("Date Header was not specified
↪");
            throw new WebApplicationException(cause, Response.Status.BAD_REQUEST);
        }
    }
}
```

This example filter checks the request for the “Date” header, and denies the request if it was missing. Otherwise, the request is passed through.

Filters can be dynamically bound to resource methods using `DynamicFeature`:

```
@Provider
public class DateRequiredFeature implements DynamicFeature {
    @Override
    public void configure(ResourceInfo resourceInfo, FeatureContext context) {
        if (resourceInfo.getResourceMethod().getAnnotation(DateRequired.class) != null) {
            context.register(DateNotSpecifiedFilter.class);
        }
    }
}
```

The `DynamicFeature` is invoked by the Jersey runtime when the application is started. In this example, the feature checks for methods that are annotated with `@DateRequired` and registers the `DateNotSpecified` filter on those methods only.

You typically register the feature in your Application class, like so:

```
environment.jersey().register(DateRequiredFeature.class);
```

4.17.2 Servlet filters

Another way to create filters is by creating servlet filters. They offer a way to register filters that apply both to servlet requests as well as resource requests. Jetty comes with a few `bundled` filters which may already suit your needs. If you want to create your own filter, this example demonstrates a servlet filter analogous to the previous example:

```
public class DateNotSpecifiedServletFilter implements jakarta.servlet.Filter {
    // Other methods in interface omitted for brevity

    @Override
    public void doFilter(ServletRequest request, ServletResponse response, FilterChain
    ↪chain) throws IOException, ServletException {
        if (request instanceof HttpServletRequest) {
            String dateHeader = ((HttpServletRequest) request).getHeader(HttpHeaders.
            ↪DATE);

            if (dateHeader != null) {
                chain.doFilter(request, response); // This signals that the request
            ↪should pass this filter
            } else {
                HttpServletResponse httpResponse = (HttpServletResponse) response;
                httpResponse.setStatus(HttpStatus.BAD_REQUEST_400);
                httpResponse.getWriter().print("Date Header was not specified");
            }
        }
    }
}
```

This servlet filter can then be registered in your Application class by wrapping it in `FilterHolder` and adding it to the application context together with a specification for which paths this filter will be active. Here’s an example:

```
environment.servlets().addFilter("DateNotSpecifiedServletFilter", new  
↳DateNotSpecifiedServletFilter()  
        .addMappingForUrlPatterns(EnumSet.of(DispatcherType.REQUEST), true,  
↳ "/"*");
```

4.18 How it's glued together

When your application starts up, it will spin up a Jetty HTTP server, see `DefaultServerFactory`. This server will have two handlers, one for your application port and the other for your admin port. The admin handler creates and registers the `AdminServlet`. This has a handle to all of the application healthchecks and metrics via the `ServletContext`.

The application port has an `HttpServlet` as well, this is composed of `DropwizardResourceConfig`, which is an extension of Jersey's resource configuration that performs scanning to find root resource and provider classes. Ultimately when you call `env.jersey().register(new SomeResource())`, you are adding to the `DropwizardResourceConfig`. This config is a jersey `Application`, so all of your application resources are served from one `Servlet`

`DropwizardResourceConfig` is where the various `ResourceMethodDispatchAdapter` are registered to enable the following functionality:

- Resource method requests with `@Timed`, `@Metered`, `@ExceptionMetered` are delegated to special dispatchers which decorate the metric telemetry
- Resources that return `Optional` are unboxed. Present returns underlying type, and non-present 404s
- Resource methods that are annotated with `@CacheControl` are delegated to a special dispatcher that decorates on the cache control headers
- Enables using Jackson to parse request entities into objects and generate response entities from objects, all while performing validation

DROPWIZARD DEPENDENCY INJECTION

Dropwizard provides you with simple dependency injection mechanism, using HK2, out-of-the-box, and you can add support for more advanced DI by using Guice bundle.

5.1 Dependency Injection Using HK2

The underlying library for out-of-the-box dependency injection mechanism in Dropwizard is Eclipse's [HK2](#), a CDI-compliant dependency injection framework.

To create a dependency injection configuration that can be overridden during test execution for mocking purposes, put it into your app Configuration for bundle to consume:

```
public interface DependencyInjectionConfiguration {
    List<Class<?>> getSingletons();
    List<NamedProperty<? extends Object>> getNamedProperties();
}

public class NamedProperty<T> {
    private final String id;
    private final T value;
    private final Class<T> clazz;

    @JsonCreator
    public NamedProperty(@JsonProperty("id") String id, @JsonProperty("value") T value,
↳ @JsonProperty("clazz") Class<T> clazz) {
        this.id = id;
        this.value = value;
        this.clazz = clazz;
    }

    public String getId() {
        return id;
    }

    public T getValue() {
        return value;
    }

    public Class<T> getClazz() {
        return clazz;
    }
}
```

(continues on next page)

(continued from previous page)

```

}

public class ExampleConfiguration extends Configuration implements
↳DependencyInjectionConfiguration {

    protected Class<?> getUserRepository() {
        return UserRepository.class;
    }

    @Override
    public List<Class<?>> getSingletons() {
        final List<Class<?>> result = new ArrayList();
        result.add(getUserRepository());
        result.add(UserResource.class);

        return result;
    }

    @Override
    public List<NamedProperty<? extends Object>> getNamedProperties() {
        final List<NamedProperty<? extends Object>> result = new ArrayList<>();
        result.add(new NamedProperty<>("dbUser", "dummy_db_user", String.class));

        return result;
    }
}

```

Then implement a bundle for DI:

```

public class DependencyInjectionBundle implements ConfiguredBundle
↳<DependencyInjectionConfiguration> {

    @Override
    public void run(DependencyInjectionConfiguration configuration, Environment
↳environment) throws Exception {
        environment
            .jersey()
            .register(
                new AbstractBinder() {
                    @Override
                    protected void configure() {
                        for (Class<?> singletonClass : configuration.
↳getSingletons()) {
                            bindAsContract(singletonClass).in(Singleton.class);
                        }

                        for (NamedProperty<? extends Object> namedProperty :
↳configuration.getNamedProperties()) {
                            bind((Object) namedProperty.getValue()).to((Class<Object>
↳) namedProperty.getClazz()).named(namedProperty.getId());
                        }
                    }
                }
            )
    }
}

```

(continues on next page)

(continued from previous page)

```

        }
    };
}

```

Then, in your application's run method, create a new `DependencyInjectionBundle`:

```

@Override
public void run(ExampleConfiguration config,
               Environment environment) {
    final DependencyInjectionBundle dependencyInjectionBundle = new
↳DependencyInjectionBundle();
    dependencyInjectionBundle.run(configuration, environment);
}

```

This allows you to use CDI annotations to control your dependency injection:

```

@Singleton
public class UserResource {
    private final UserRepository userRepository;

    @Inject
    public UserResource(UserRepository userRepository) {
        this.userRepository = userRepository;
    }
}

@Singleton
public class UserRepository {
    private final String dbUser;

    @Inject
    public UserRepository(@Named("dbUser") String dbUser) {
        this.dbUser = dbUser;
    }
}

```

Then you can provide alternate configuration for testing purposes:

```

public class TestConfiguration extends ExampleConfiguration {

    @Override
    protected Class<?> getUserRepository() {
        return MockUserRepository.class;
    }
}

@DisplayName("User endpoint")
@ExtendWith(DropwizardExtensionsSupport.class)
public class UserControllerTests {
    public static final DropwizardAppExtension<TestConfiguration> app = new
↳DropwizardAppExtension<>(ExampleApplication.class, new TestConfiguration());
}

```

Note: the `@Singleton` annotation is only effective for Dropwizard resources. For custom classes, don't forget to register them as shown above with `bindAsContract(singletonClass).in(Singleton.class)`.

DROPWIZARD CLIENT

The `dropwizard-client` module provides you with two different performant, instrumented HTTP clients so you can integrate your service with other web services: Apache `HttpClient` and Jersey Client.

```
<dependency>
  <groupId>io.dropwizard</groupId>
  <artifactId>dropwizard-client</artifactId>
</dependency>
```

6.1 Apache HttpClient

The underlying library for `dropwizard-client` is Apache's `HttpClient`, a full-featured, well-tested HTTP client library.

To create a *managed*, instrumented `HttpClient` instance, your *configuration class* needs an *http client configuration* instance:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private HttpClientConfiguration httpClient = new HttpClientConfiguration();

    @JsonProperty("httpClient")
    public HttpClientConfiguration getHttpClientConfiguration() {
        return httpClient;
    }

    @JsonProperty("httpClient")
    public void setHttpClientConfiguration(HttpClientConfiguration httpClient) {
        this.httpClient = httpClient;
    }
}
```

Then, in your application's `run` method, create a new `HttpClientBuilder`:

```
@Override
public void run(ExampleConfiguration config,
               Environment environment) {
    final HttpClient httpClient = new HttpClientBuilder(environment).using(config.
```

(continues on next page)

(continued from previous page)

```
↪getHttpClientConfiguration()  
    .build(getName());  
    environment.jersey().register(new ExternalServiceResource(httpClient));  
}
```

6.1.1 Metrics

Dropwizard's `HttpClientBuilder` actually gives you an instrumented subclass which tracks the following pieces of data:

`org.apache.http.conn.ClientConnectionManager.available-connections`

The number of idle connections ready to be used to execute requests.

`org.apache.http.conn.ClientConnectionManager.leased-connections`

The number of persistent connections currently being used to execute requests.

`org.apache.http.conn.ClientConnectionManager.max-connections`

The maximum number of allowed connections.

`org.apache.http.conn.ClientConnectionManager.pending-connections`

The number of connection requests being blocked awaiting a free connection.

`org.apache.http.client.HttpClient.get-requests`

The rate at which GET requests are being sent.

`org.apache.http.client.HttpClient.post-requests`

The rate at which POST requests are being sent.

`org.apache.http.client.HttpClient.head-requests`

The rate at which HEAD requests are being sent.

`org.apache.http.client.HttpClient.put-requests`

The rate at which PUT requests are being sent.

`org.apache.http.client.HttpClient.delete-requests`

The rate at which DELETE requests are being sent.

`org.apache.http.client.HttpClient.options-requests`

The rate at which OPTIONS requests are being sent.

`org.apache.http.client.HttpClient.trace-requests`

The rate at which TRACE requests are being sent.

`org.apache.http.client.HttpClient.connect-requests`

The rate at which CONNECT requests are being sent.

`org.apache.http.client.HttpClient.move-requests`

The rate at which MOVE requests are being sent.

`org.apache.http.client.HttpClient.patch-requests`

The rate at which PATCH requests are being sent.

`org.apache.http.client.HttpClient.other-requests`

The rate at which requests with none of the above methods are being sent.

Note: The naming strategy for the metrics associated requests is configurable. Specifically, the last part e.g. `get-requests`. What is displayed is `HttpClientMetricNameStrategies.METHOD_ONLY`, you can also in-

clude the host via `HttpClientMetricNameStrategies.HOST_AND_METHOD` or a url without query string via `HttpClientMetricNameStrategies.QUERYLESS_URL_AND_METHOD`

6.2 Jersey Client

If `HttpClient` is too low-level for you, Dropwizard also supports Jersey's `Client API`. Jersey's `Client` allows you to use all of the server-side media type support that your service uses to, for example, deserialize application/json request entities as POJOs.

To create a *managed*, instrumented `JerseyClient` instance, your *configuration class* needs an *jersey client configuration* instance:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private JerseyClientConfiguration jerseyClient = new JerseyClientConfiguration();

    @JsonProperty("jerseyClient")
    public JerseyClientConfiguration getJerseyClientConfiguration() {
        return jerseyClient;
    }

    @JsonProperty("jerseyClient")
    public void setJerseyClientConfiguration(JerseyClientConfiguration jerseyClient) {
        this.jerseyClient = jerseyClient;
    }
}
```

Then, in your service's run method, create a new `JerseyClientBuilder`:

```
@Override
public void run(ExampleConfiguration config,
               Environment environment) {

    final Client client = new JerseyClientBuilder(environment).using(config.
↳getJerseyClientConfiguration())
                                                                .build(getName());
    environment.jersey().register(new ExternalServiceResource(client));
}
```

6.2.1 Configuration

The `Client` that Dropwizard creates deviates from the *Jersey Client Configuration* defaults. The default, in Jersey, is for a client to never timeout reading or connecting in a request, while in Dropwizard, the default is 500 milliseconds.

There are a couple of ways to change this behavior. The recommended way is to modify the *YAML configuration*. Alternatively, set the properties on the `JerseyClientConfiguration`, which will take effect for all built clients. On a per client basis, the configuration can be changed by utilizing the `property` method and, in this case, the *Jersey Client Properties* can be used.

Warning: Do not try to change Jersey properties using [Jersey Client Properties](#) through the `withProperty(String propertyName, Object propertyValue)` method on the `JerseyClientBuilder`, because by default it's configured by Dropwizard's `HttpClientBuilder`, so the Jersey properties are ignored.

6.2.2 Rx Usage

To increase the ergonomics of asynchronous client requests, Jersey allows creation of `rx-clients`. You can instruct Dropwizard to create such a client (RxJava2):

```
@Override
public void run(ExampleConfiguration config,
                Environment environment) {

    final Client client =
        new JerseyClientBuilder(environment)
            .using(config.getJerseyClientConfiguration())
            .buildRx(getName(), RxFlowableInvokerProvider.class);
    //Any custom Service Resource that waits for Client in constructor
    environment.jersey().register(new ExternalServiceResource(client));
}
```

`RxFlowableInvokerProvider.class` is the JavaRx implementation and can be added to the pom:

```
<dependency>
  <groupId>org.glassfish.jersey.ext.rx</groupId>
  <artifactId>jersey-rx-client-rxjava2</artifactId>
</dependency>
```

Alternatively, there are RxJava, Guava, and JSR-166e implementations.

By allowing Dropwizard to create the `rx-client`, the same thread pool that is utilized by traditional synchronous and asynchronous requests, is used for `rx` requests.

6.2.3 Proxy Authentication

The client can utilise a forward proxy, supporting both Basic and NTLM authentication schemes. Basic Auth against a proxy is simple:

```
proxy:
  host: '192.168.52.11'
  port: 8080
  scheme : 'https'
  auth:
    username: 'secret'
    password: 'stuff'
  nonProxyHosts:
    - 'localhost'
    - '192.168.52.*'
    - '*.example.com'
```

NTLM Auth is configured by setting the relevant windows properties.

```
proxy:
  host: '192.168.52.11'
  port: 8080
  scheme : 'https'
  auth:
    username: 'secret'
    password: 'stuff'
    authScheme: 'NTLM'
    realm: 'realm' # optional, defaults to ANY_REALM
    hostname: 'workstation' # optional, defaults to null but may be
↔required depending on your AD environment
    domain: 'HYPERCOMPUGLOBALMEGANET' # optional, defaults to null but may be
↔required depending on your AD environment
    credentialType: 'NT'
  nonProxyHosts:
    - 'localhost'
    - '192.168.52.*'
    - '*.example.com'
```


DROPWIZARD JDBC3

The `dropwizard-jdbi3` module provides you with managed access to JDBI, a flexible and modular library for interacting with relational databases via SQL.

7.1 Configuration

To create a *managed*, instrumented Jdbi instance, your *configuration class* needs a DataSourceFactory instance:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private DataSourceFactory database = new DataSourceFactory();

    @JsonProperty("database")
    public void setDataSourceFactory(DataSourceFactory factory) {
        this.database = factory;
    }

    @JsonProperty("database")
    public DataSourceFactory getDataSourceFactory() {
        return database;
    }
}
```

Then, in your service's run method, create a new JdbiFactory:

```
@Override
public void run(ExampleConfiguration config, Environment environment) {
    final JdbiFactory factory = new JdbiFactory();
    final Jdbi jdbi = factory.build(environment, config.getDataSourceFactory(),
    ↪ "postgresql");
    environment.jersey().register(new UserResource(jdbi));
}
```

This will create a new *managed* connection pool to the database, a *health check* for connectivity to the database, and a new Jdbi instance for you to use.

Your service's configuration file will then look like this:

```
database:
  # the name of your JDBC driver
```

(continues on next page)

```
driverClass: org.postgresql.Driver

# the username
user: pg-user

# the password
password: iAMs00perSecrEET

# the JDBC URL
url: jdbc:postgresql://db.example.com/db-prod

# any properties specific to your JDBC driver:
properties:
  charSet: UTF-8

# the maximum amount of time to wait on an empty pool before throwing an exception
maxWaitForConnection: 1s

# the SQL query to run when validating a connection's liveness
validationQuery: "/* MyService Health Check */ SELECT 1"

# the timeout before a connection validation queries fail
validationQueryTimeout: 3s

# the minimum number of connections to keep open
minSize: 8

# the maximum number of connections to keep open
maxSize: 32

# whether or not idle connections should be validated
checkConnectionWhileIdle: false

# the amount of time to sleep between runs of the idle connection validation,
↳ abandoned cleaner and idle pool resizing
evictionInterval: 10s

# the minimum amount of time an connection must sit idle in the pool before it is
↳ eligible for eviction
minIdleTime: 1 minute
```

7.2 Plugins

JDBI3 is built using [plugins](#) to add features to its core implementation. Dropwizard adds the [sqlobject](#) and [guava](#) plugins by default, but you are free to add other existing plugins you might need or create your own.

7.3 Usage

We highly recommend you use JDBI's [SQL Objects API](#), which allows you to write DAO classes as interfaces:

```
public interface MyDAO {
    @SqlUpdate("create table something (id int primary key, name varchar(100))")
    void createSomethingTable();

    @SqlUpdate("insert into something (id, name) values (:id, :name)")
    void insert(@Bind("id") int id, @Bind("name") String name);

    @SqlQuery("select name from something where id = :id")
    String findNameById(@Bind("id") int id);
}

final MyDAO dao = database.onDemand(MyDAO.class);
```

This ensures your DAO classes are trivially mockable, as well as encouraging you to extract mapping code (e.g., `RowMapper` -> domain objects) into testable, reusable classes.

7.4 Exception Handling

By adding the `JdbiExceptionsBundle` to your *application*, Dropwizard will automatically unwrap any thrown `SQLException` or `JdbiException` instances. This is critical for debugging, since otherwise only the common wrapper exception's stack trace is logged.

7.5 Prepend Comments

If you're using JDBI's [SQL Objects API](#) (and you should be), `dropwizard-jdbi3` will automatically prepend the SQL object's class and method name to the SQL query as an SQL comment:

```
/* com.example.service.dao.UserDAO.findByName */
SELECT id, name, email
FROM users
WHERE name = 'Coda';
```

This will allow you to quickly determine the origin of any slow or misbehaving queries.

DROPWIZARD MIGRATIONS

The `dropwizard-migrations` module provides you with a wrapper for Liquibase database refactoring.

8.1 Configuration

Like *Dropwizard JDBI3*, your *configuration class* needs a `DataSourceFactory` instance:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private DataSourceFactory database = new DataSourceFactory();

    @JsonProperty("database")
    public DataSourceFactory getDataSourceFactory() {
        return database;
    }
}
```

8.2 Adding The Bundle

Then, in your application's `initialize` method, add a new `MigrationsBundle` subclass:

```
@Override
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.addBundle(new MigrationsBundle<ExampleConfiguration>() {
        @Override
        public DataSourceFactory getDataSourceFactory(ExampleConfiguration
↪ configuration) {
            return configuration.getDataSourceFactory();
        }
    });
}
```

If you are using *Dropwizard Hibernate* or *Dropwizard JDBI3* in your application, you can use these techniques within a `CustomChange` to make bigger data migrations. Therefore you need to provide an instance of these to the `MigrationsBundle` like in the following example:

```
@Override
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
```

(continues on next page)

(continued from previous page)

```

bootstrap.addBundle(new MigrationsBundle<ExampleConfiguration>() {
    @Override
    public DataSourceFactory getDataSourceFactory(ExampleConfiguration
↳ configuration) {
        return configuration.getDataSourceFactory();
    }

    @Override
    public Map<String, Object> getScopedObjects() {
        Map<String, Object> scopedObjects = new HashMap<>();
        scopedObjects.put("hibernateSessionFactory", hibernateBundle
↳ getSessionFactory());
        return scopedObjects;
    }
});
}

```

In your CustomChange you can retrieve the registered SessionFactory with this code:

```

public void execute(Database database) throws CustomChangeException {
    Scope.getCurrentScope().get("hibernateSessionFactory", SessionFactory.class);
}

```

8.3 Defining Migrations

Your database migrations are stored in your Dropwizard project, in `src/main/resources/migrations.xml`. This file will be packaged with your application, allowing you to run migrations using your application's command-line interface. You can change the name of the migrations file used by overriding the `getMigrationsFileName()` method in `MigrationsBundle`.

For example, to create a new people table, you might create an initial `migrations.xml` like this:

```

<?xml version="1.0" encoding="UTF-8"?>

<databaseChangeLog
    xmlns="http://www.liquibase.org/xml/ns/dbchangelog"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.liquibase.org/xml/ns/dbchangelog
    http://www.liquibase.org/xml/ns/dbchangelog/dbchangelog-3.1.xsd">

    <changeSet id="1" author="codahale">
        <createTable tableName="people">
            <column name="id" type="bigint" autoIncrement="true">
                <constraints primaryKey="true" nullable="false"/>
            </column>
            <column name="fullName" type="varchar(255)">
                <constraints nullable="false"/>
            </column>
            <column name="jobTitle" type="varchar(255)"/>
        </createTable>
    </changeSet>

```

(continues on next page)

(continued from previous page)

```
</changeSet>  
</databaseChangeLog>
```

For more information on available database refactorings, check the [Liquibase](#) documentation.

8.4 Checking Your Database's State

To check the state of your database, use the `db status` command:

```
java -jar hello-world.jar db status helloworld.yml
```

8.5 Dumping Your Schema

If your database already has an existing schema and you'd like to pre-seed your `migrations.xml` document, you can run the `db dump` command:

```
java -jar hello-world.jar db dump helloworld.yml
```

This will output a [Liquibase](#) change log with a changeset capable of recreating your database.

8.6 Tagging Your Schema

To tag your schema at a particular point in time (e.g., to make rolling back easier), use the `db tag` command:

```
java -jar hello-world.jar db tag helloworld.yml 2012-10-08-pre-user-move
```

8.7 Migrating Your Schema

To apply pending changesets to your database schema, run the `db migrate` command:

```
java -jar hello-world.jar db migrate helloworld.yml
```

Warning: This will potentially make irreversible changes to your database. Always check the pending DDL scripts by using the `--dry-run` flag first. This will output the SQL to be run to stdout.

Note: To apply only a specific number of pending changesets, use the `--count` flag.

8.8 Rolling Back Your Schema

To roll back changesets which have already been applied, run the `db rollback` command. You will need to specify either a **tag**, a **date**, or a **number of changesets** to roll back to:

```
java -jar hello-world.jar db rollback helloworld.yml --tag 2012-10-08-pre-user-move
```

Warning: This will potentially make irreversible changes to your database. Always check the pending DDL scripts by using the `--dry-run` flag first. This will output the SQL to be run to stdout.

8.9 Testing Migrations

To verify that a set of pending changesets can be fully rolled back, use the `db test` command, which will migrate forward, roll back to the original state, then migrate forward again:

```
java -jar hello-world.jar db test helloworld.yml
```

Warning: Do not run this in production, for obvious reasons.

8.10 Preparing A Rollback Script

To prepare a rollback script for pending changesets *before* they have been applied, use the `db prepare-rollback` command:

```
java -jar hello-world.jar db prepare-rollback helloworld.yml
```

This will output a DDL script to stdout capable of rolling back all unapplied changesets.

8.11 Generating Documentation

To generate HTML documentation on the current status of the database, use the `db generate-docs` command:

```
java -jar hello-world.jar db generate-docs helloworld.yml ~/db-docs/
```

8.12 Dropping All Objects

To drop all objects in the database, use the `db drop-all` command:

```
java -jar hello-world.jar db drop-all --confirm-delete-everything helloworld.yml
```

Warning: You need to specify the `--confirm-delete-everything` flag because this command **deletes everything in the database**. Be sure you want to do that first.

8.13 Fast-Forwarding Through A Changeset

To mark a pending changeset as applied (e.g., after having backfilled your `migrations.xml` with `db dump`), use the `db fast-forward` command:

```
java -jar hello-world.jar db fast-forward helloworld.yml
```

This will mark the next pending changeset as applied. You can also use the `--all` flag to mark all pending changesets as applied.

8.14 Support For Adding Multiple Migration Bundles

Assuming migrations need to be done for two different databases, you would need to have two different data source factories:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private DataSourceFactory database1 = new DataSourceFactory();

    @Valid
    @NotNull
    private DataSourceFactory database2 = new DataSourceFactory();

    @JsonProperty("database1")
    public DataSourceFactory getDb1DataSourceFactory() {
        return database1;
    }

    @JsonProperty("database2")
    public DataSourceFactory getDb2DataSourceFactory() {
        return database2;
    }
}
```

Now multiple migration bundles can be added with unique names like so:

```
@Override
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.addBundle(new MigrationsBundle<ExampleConfiguration>() {
        @Override
        public DataSourceFactory getDataSourceFactory(ExampleConfiguration
↳ configuration) {
            return configuration.getDb1DataSourceFactory();
        }

        @Override
        public String name() {
            return "db1";
        }
    });
}
```

(continues on next page)

(continued from previous page)

```
bootstrap.addBundle(new MigrationsBundle<ExampleConfiguration>() {
    @Override
    public DataSourceFactory getDataSourceFactory(ExampleConfiguration_
↵configuration) {
        return configuration.getDb2DataSourceFactory();
    }

    @Override
    public String name() {
        return "db2";
    }
});
}
```

To migrate your schema:

```
java -jar hello-world.jar db1 migrate helloworld.yml
```

and

```
java -jar hello-world.jar db2 migrate helloworld.yml
```

Note: Whenever a name is added to a migration bundle, it becomes the command that needs to be run at the command line. eg: To check the state of your database, use the status command:

```
java -jar hello-world.jar db1 status helloworld.yml
```

or

```
java -jar hello-world.jar db2 status helloworld.yml
```

By default the migration bundle uses the “db” command. By overriding you can customize it to provide any name you want and have multiple migration bundles. Wherever the “db” command was being used, this custom name can be used.

There will also be a need to provide different change log migration files as well. This can be done as

```
java -jar hello-world.jar db1 migrate helloworld.yml --migrations <path_to_db1_
↵migrations.xml>
```

```
java -jar hello-world.jar db2 migrate helloworld.yml --migrations <path_to_db2_
↵migrations.xml>
```

8.15 More Information

If you are using databases supporting multiple schemas like PostgreSQL, Oracle, or H2, you can use the optional `--catalog` and `--schema` arguments to specify the database catalog and schema used for the Liquibase commands.

For more information on available commands, either use the `db --help` command, or for more detailed help on a specific command, use `db <cmd> --help`.

DROPWIZARD HIBERNATE

The `dropwizard-hibernate` module provides you with managed access to Hibernate, a powerful, industry-standard object-relation mapper (ORM).

9.1 Configuration

To create a *managed*, instrumented `SessionFactory` instance, your *configuration class* needs a `DataSourceFactory` instance:

```
public class ExampleConfiguration extends Configuration {
    @Valid
    @NotNull
    private DataSourceFactory database = new DataSourceFactory();

    @JsonProperty("database")
    public DataSourceFactory getDataSourceFactory() {
        return database;
    }

    @JsonProperty("database")
    public void setDataSourceFactory(DataSourceFactory dataSourceFactory) {
        this.database = dataSourceFactory;
    }
}
```

Then, add a `HibernateBundle` instance to your application class, specifying your entity classes and how to get a `DataSourceFactory` from your configuration subclass:

```
private final HibernateBundle<ExampleConfiguration> hibernate = new HibernateBundle
    <ExampleConfiguration>(Person.class) {
    @Override
    public DataSourceFactory getDataSourceFactory(ExampleConfiguration configuration) {
        return configuration.getDataSourceFactory();
    }
};

@Override
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.addBundle(hibernate);
}
```

(continues on next page)

(continued from previous page)

```
@Override
public void run(ExampleConfiguration config, Environment environment) {
    final PersonDAO dao = new PersonDAO(hibernate.getSessionFactory());
    environment.jersey().register(new UserResource(dao));
}
```

This will create a new *managed* connection pool to the database, a *health check* for connectivity to the database, and a new `SessionFactory` instance for you to use in your DAO classes.

Your application's configuration file will then look like this:

```
database:
# the name of your JDBC driver
driverClass: org.postgresql.Driver

# the username
user: pg-user

# the password
password: iAMs00perSecrEET

# the JDBC URL
url: jdbc:postgresql://db.example.com/db-prod

# any properties specific to your JDBC driver:
properties:
    charset: UTF-8
    hibernate.dialect: org.hibernate.dialect.PostgreSQLDialect

# the maximum amount of time to wait on an empty pool before throwing an exception
maxWaitForConnection: 1s

# the SQL query to run when validating a connection's liveness
validationQuery: "/* MyApplication Health Check */ SELECT 1"

# the minimum number of connections to keep open
minSize: 8

# the maximum number of connections to keep open
maxSize: 32

# whether or not idle connections should be validated
checkConnectionWhileIdle: false
```

9.2 Usage

9.2.1 Data Access Objects

Dropwizard comes with `AbstractDAO`, a minimal template for entity-specific DAO classes. It contains type-safe wrappers for `SessionFactory`'s common operations:

```
public class PersonDAO extends AbstractDAO<Person> {
    public PersonDAO(SessionFactory factory) {
        super(factory);
    }

    public Person findById(Long id) {
        return get(id);
    }

    public long create(Person person) {
        return persist(person).getId();
    }

    public List<Person> findAll() {
        return list(namedTypedQuery("com.example.helloworld.core.Person.findAll"));
    }
}
```

9.2.2 Transactional Resource Methods

Dropwizard uses a declarative method of scoping transactional boundaries. Not all resource methods actually require database access, so the `@UnitOfWork` annotation is provided:

```
@GET
@Path("/{id}")
@Timed
@UnitOfWork
public Person findPerson(@PathParam("id") LongParam id) {
    return dao.findById(id.get());
}
```

This will automatically open a session, begin a transaction, call `findById`, commit the transaction, and finally close the session. If an exception is thrown, the transaction is rolled back.

If you are using more than one Hibernate bundle in your application, you can repeat the `@UnitOfWork` annotation:

```
@GET
@Path("...")
@Timed
@UnitOfWork(value = "hibernate.<db-name-1>")
@UnitOfWork(value = "hibernate.<db-name-2>")
public ...
```

In this case, the above specified behaviour applies to all given databases.

If multiple `@UnitOfWork` annotations with the same `value()` are provided, the last one is used.

Important: The Hibernate session is closed **before** your resource method's return value (e.g., the `Person` from the database), which means your resource method (or DAO) is responsible for initializing all lazily-loaded collections, etc., before returning. Otherwise, you'll get a `LazyInitializationException` thrown in your template (or null values produced by Jackson).

9.2.3 Transactional Resource Methods Outside Jersey Resources

Currently creating transactions with the `@UnitOfWork` annotation works out-of-box only for resources managed by Jersey. If you want to use it outside Jersey resources, e.g. in authenticators, you should instantiate your class with `UnitOfWorkAwareProxyFactory`.

```
SessionDao dao = new SessionDao(hibernateBundle.getSessionFactory());
ExampleAuthenticator exampleAuthenticator = new
↳UnitOfWorkAwareProxyFactory(hibernateBundle)
    .create(ExampleAuthenticator.class, SessionDao.class, dao);
```

It will create a proxy of your class, which will open a Hibernate session with a transaction around methods with the `@UnitOfWork` annotation.

9.3 Prepended Comments

Dropwizard automatically configures Hibernate to prepend a comment describing the context of all queries:

```
/* load com.example.helloworld.core.Person */
select
  person0_.id as id0_0_,
  person0_.fullName as fullName0_0_,
  person0_.jobTitle as jobTitle0_0_
from people person0_
where person0_.id=?
```

This will allow you to quickly determine the origin of any slow or misbehaving queries.

DROPWIZARD AUTHENTICATION

The `dropwizard-auth` client provides authentication using either HTTP Basic Authentication or OAuth2 bearer tokens.

10.1 Authenticators

An authenticator is a strategy class which, given a set of client-provided credentials, possibly returns a principal (i.e., the person or entity on behalf of whom your service will do something).

Authenticators implement the `Authenticator<C, P extends Principal>` interface, which has a single method:

```
public class ExampleAuthenticator implements Authenticator<BasicCredentials, User> {
    @Override
    public Optional<User> authenticate(BasicCredentials credentials) throws
↳ AuthenticationException {
        if ("secret".equals(credentials.getPassword())) {
            return Optional.of(new User(credentials.getUsername()));
        }
        return Optional.empty();
    }
}
```

This authenticator takes *basic auth credentials* and if the client-provided password is `secret`, authenticates the client as a `User` with the client-provided username.

If the password doesn't match, an absent `Optional` is returned instead, indicating that the credentials are invalid.

Warning: It's important for authentication services not to provide too much information in their errors. The fact that a username or email has an account may be meaningful to an attacker, so the `Authenticator` interface doesn't allow you to distinguish between a bad username and a bad password. You should only throw an `AuthenticationException` if the authenticator is **unable** to check the credentials (e.g., your database is down).

10.1.1 Caching

Because the backing data stores for authenticators may not handle high throughput (an RDBMS or LDAP server, for example), Dropwizard provides a decorator class which provides caching:

```
SimpleAuthenticator simpleAuthenticator = new SimpleAuthenticator();
CachingAuthenticator<BasicCredentials, User> cachingAuthenticator = new
↳CachingAuthenticator<>(
    metricRegistry, simpleAuthenticator,
    config.getAuthenticationCachePolicy());
```

Dropwizard can parse Caffeine's CaffeineSpec from the configuration policy, allowing your configuration file to look like this:

```
authenticationCachePolicy: maximumSize=10000, expireAfterAccess=10m
```

This caches up to 10,000 principals, evicting stale entries after 10 minutes.

10.2 Authorizer

An authorizer is a strategy class which, given a principal and a role, decides if access is granted to the principal.

The authorizer implements the `Authorizer<P extends Principal>` interface, which has a single method:

```
public class ExampleAuthorizer implements Authorizer<User> {
    @Override
    public boolean authorize(User user, String role) {
        return user.getName().equals("good-guy") && role.equals("ADMIN");
    }
}
```

10.3 Basic Authentication

The `AuthDynamicFeature` with the `BasicCredentialAuthFilter` and `RolesAllowedDynamicFeature` enables HTTP Basic authentication and authorization; requires an authenticator which takes instances of `BasicCredentials`. If you don't use authorization, then `RolesAllowedDynamicFeature` is not required.

```
@Override
public void run(ExampleConfiguration configuration,
    Environment environment) {
    environment.jersey().register(new AuthDynamicFeature(
        new BasicCredentialAuthFilter.Builder<User>()
            .setAuthenticator(new ExampleAuthenticator())
            .setAuthorizer(new ExampleAuthorizer())
            .setRealm("SUPER SECRET STUFF")
            .buildAuthFilter());
    environment.jersey().register(RolesAllowedDynamicFeature.class);
    //If you want to use @Auth to inject a custom Principal type into your resource
    environment.jersey().register(new AuthValueFactoryProvider.Binder<>(User.class));
}
```

10.4 OAuth2

The `AuthDynamicFeature` with `OAuthCredentialAuthFilter` and `RolesAllowedDynamicFeature` enables OAuth2 bearer-token authentication and authorization; requires an authenticator which takes instances of `String`. If you don't use authorization, then `RolesAllowedDynamicFeature` is not required.

```
@Override
public void run(ExampleConfiguration configuration,
                Environment environment) {
    environment.jersey().register(new AuthDynamicFeature(
        new OAuthCredentialAuthFilter.Builder<User>()
            .setAuthenticator(new ExampleOAuthAuthenticator())
            .setAuthorizer(new ExampleAuthorizer())
            .setPrefix("Bearer")
            .buildAuthFilter()));
    environment.jersey().register(RolesAllowedDynamicFeature.class);
    //If you want to use @Auth to inject a custom Principal type into your resource
    environment.jersey().register(new AuthValueFactoryProvider.Binder<>(User.class));
}
```

10.5 Chained Factories

The `ChainedAuthFilter` enables usage of various authentication factories at the same time.

```
@Override
public void run(ExampleConfiguration configuration,
                Environment environment) {
    AuthFilter basicCredentialAuthFilter = new BasicCredentialAuthFilter.Builder<>()
        .setAuthenticator(new ExampleBasicAuthenticator())
        .setAuthorizer(new ExampleAuthorizer())
        .setPrefix("Basic")
        .buildAuthFilter();

    AuthFilter oauthCredentialAuthFilter = new OAuthCredentialAuthFilter.Builder<>()
        .setAuthenticator(new ExampleOAuthAuthenticator())
        .setAuthorizer(new ExampleAuthorizer())
        .setPrefix("Bearer")
        .buildAuthFilter();

    List<AuthFilter> filters = Lists.newArrayList(basicCredentialAuthFilter,
↪oauthCredentialAuthFilter);
    environment.jersey().register(new AuthDynamicFeature(new
↪ChainedAuthFilter(filters));
    environment.jersey().register(RolesAllowedDynamicFeature.class);
    //If you want to use @Auth to inject a custom Principal type into your resource
    environment.jersey().register(new AuthValueFactoryProvider.Binder<>(User.class));
}
```

For this to work properly, all chained factories must produce the same type of principal, here `User`.

10.6 Protecting Resources

There are two ways to protect a resource. You can mark your resource method with one of the following annotations:

- `@PermitAll`. All authenticated users will have access to the method.
- `@RolesAllowed`. Access will be granted to the users with the specified roles.
- `@DenyAll`. No access will be granted to anyone.

Note: You can use `@RolesAllowed`, `@PermitAll` on the class level. Method annotations take precedence over the class ones.

Alternatively, you can annotate the parameter representing your principal with `@Auth`. Note you must register a jersey provider to make this work.

```
environment.jersey().register(new AuthValueFactoryProvider.Binder<>(User.class));

@RolesAllowed("ADMIN")
@GET
public SecretPlan getSecretPlan(@Auth User user) {
    return dao.findPlanForUser(user);
}
```

You can also access the Principal by adding a parameter to your method `@Context SecurityContext context`. Note this will not automatically register the servlet filter which performs authentication. You will still need to add one of `@PermitAll`, `@RolesAllowed`, or `@DenyAll`. This is not the case with `@Auth`. When that is present, the auth filter is automatically registered to facilitate users upgrading from older versions of Dropwizard

```
@RolesAllowed("ADMIN")
@GET
public SecretPlan getSecretPlan(@Context SecurityContext context) {
    User userPrincipal = (User) context.getUserPrincipal();
    return dao.findPlanForUser(user);
}
```

If there are no provided credentials for the request, or if the credentials are invalid, the provider will return a scheme-appropriate 401 Unauthorized response without calling your resource method.

10.6.1 Optional protection

Resource methods can be *optionally* protected by representing the principal as an `Optional`. In such cases, the `Optional` resource method argument will be populated with the principal, if present. Otherwise, the argument will be `Optional.empty`.

For instance, say you have an endpoint that should display a logged-in user's name, but return an anonymous reply for unauthenticated requests. You need to implement a custom filter which injects a security context containing the principal if it exists, without performing authentication.

```
@GET
public String getGreeting(@Auth Optional<User> userOpt) {
    if (userOpt.isPresent()) {
        return "Hello, " + userOpt.get().getName() + "!";
    }
}
```

(continues on next page)

(continued from previous page)

```

    } else {
        return "Greetings, anonymous visitor!"
    }
}

```

For optionally-protected resources, requests with invalid auth will be treated the same as those with no provided auth credentials. That is to say, requests that `_fail_` to meet an authenticator or authorizer's requirements result in an empty principal being passed to the resource method.

10.7 Testing Protected Resources

Add this dependency into your `pom.xml` file:

```

<dependencies>
  <dependency>
    <groupId>io.dropwizard</groupId>
    <artifactId>dropwizard-testing</artifactId>
    <version>${dropwizard.version}</version>
  </dependency>
  <dependency>
    <groupId>org.glassfish.jersey.test-framework.providers</groupId>
    <artifactId>jersey-test-framework-provider-grizzly2</artifactId>
    <version>${jersey.version}</version>
    <exclusions>
      <exclusion>
        <groupId>jakarta.servlet</groupId>
        <artifactId>jakarta.servlet-api</artifactId>
      </exclusion>
      <exclusion>
        <groupId>junit</groupId>
        <artifactId>junit</artifactId>
      </exclusion>
    </exclusions>
  </dependency>
</dependencies>

```

10.7.1 OAuth Example

When you build your `ResourceExtension`, add the `GrizzlyWebTestContainerFactory` line.

```

@ExtendWith(DropwizardExtensionsSupport.class)
public class OAuthResourceTest {

    public ResourceExtension resourceExtension = ResourceExtension
        .builder()
        .setTestContainerFactory(new GrizzlyWebTestContainerFactory())
        .addProvider(new AuthDynamicFeature(new OAuthCredentialAuthFilter.Builder
    ↪ <User>()

            .setAuthenticator(new MyOAuthAuthenticator())

```

(continues on next page)

(continued from previous page)

```

        .setAuthorizer(new MyAuthorizer())
        .setRealm("SUPER SECRET STUFF")
        .setPrefix("Bearer")
        .buildAuthFilter()))
    .addProvider(RolesAllowedDynamicFeature.class)
    .addProvider(new AuthValueFactoryProvider.Binder<>(User.class))
    .addResource(new ProtectedResource())
    .build();
}

```

Note that you need to set the token header manually.

```

@Test
public void testProtected() throws Exception {
    final Response response = resourceExtension.target("/protected")
        .request(MediaType.APPLICATION_JSON_TYPE)
        .header("Authorization", "Bearer TOKEN")
        .get();

    assertThat(response.getStatus()).isEqualTo(200);
}

```

10.7.2 BasicAuth Example

When you build your ResourceExtension, add the GrizzlyWebTestContainerFactory line.

```

@ExtendWith(DropwizardExtensionsSupport.class)
public class OAuthResourceTest {
    public ResourceExtension resourceExtension = ResourceExtension
        .builder()
        .setTestContainerFactory(new GrizzlyWebTestContainerFactory())
        .addProvider(new AuthDynamicFeature(new BasicCredentialAuthFilter.Builder
↪<User>()

            .setAuthenticator(new MyBasicAuthenticator())
            .setAuthorizer(new MyBasicAuthorizer())
            .buildAuthFilter()))
        .addProvider(RolesAllowedDynamicFeature.class)
        .addProvider(new AuthValueFactoryProvider.Binder<>(User.class))
        .addResource(new ProtectedResource())
        .build()
}

```

Note that you need to set the authorization header manually.

```

@Test
public void testProtectedResource(){

    String credential = "Basic " + Base64.getEncoder().encodeToString("test@gmail.
↪com:secret".getBytes())

    Response response = resourceExtension

```

(continues on next page)

(continued from previous page)

```

        .target("/protected")
        .request()
        .header(HttpHeaders.AUTHORIZATION, credential)
        .get();

    Assert.assertEquals(200, response.getStatus());
}

```

10.8 Multiple Principals and Authenticators

In some cases you may want to use different authenticators/authentication schemes for different resources. For example you may want Basic authentication for one resource and OAuth for another resource, at the same time using a different *Principal* for each authentication scheme.

For this use case, there is the `PolymorphicAuthDynamicFeature` and the `PolymorphicAuthValueFactoryProvider`. With these two components, we can use different combinations of authentication schemes/authenticators/authorizers/principals. To use this feature, we need to do a few things:

- Register the `PolymorphicAuthDynamicFeature` with a map that maps principal types to authentication filters.
- Register the `PolymorphicAuthValueFactoryProvider` with a set of principal classes that you will be using.
- Annotate your resource method `Principal` parameters with `@Auth`.

As an example, the following code configures both OAuth and Basic authentication, using a different principal for each.

```

final AuthFilter<BasicCredentials, BasicPrincipal> basicFilter
    = new BasicCredentialAuthFilter.Builder<BasicPrincipal>()
        .setAuthenticator(new ExampleAuthenticator())
        .setRealm("SUPER SECRET STUFF")
        .buildAuthFilter();
final AuthFilter<String, OAuthPrincipal> oauthFilter
    = new OAuthCredentialAuthFilter.Builder<OAuthPrincipal>()
        .setAuthenticator(new ExampleOAuthAuthenticator())
        .setPrefix("Bearer")
        .buildAuthFilter();

final PolymorphicAuthDynamicFeature feature = new PolymorphicAuthDynamicFeature<>(
    ImmutableMap.of(
        BasicPrincipal.class, basicFilter,
        OAuthPrincipal.class, oauthFilter));
final AbstractBinder binder = new PolymorphicAuthValueFactoryProvider.Binder<>(
    ImmutableSet.of(BasicPrincipal.class, OAuthPrincipal.class));

environment.jersey().register(feature);
environment.jersey().register(binder);

```

Now we are able to do something like the following

```

@GET
public Response basicAuthResource(@Auth BasicPrincipal principal) {}

```

(continues on next page)

(continued from previous page)

```
@GET
public Response oauthResource(@Auth OAuthPrincipal principal) {}
```

The first resource method will use Basic authentication while the second one will use OAuth.

Note that with the above example, only *authentication* is configured. If you also want *authorization*, the following steps will need to be taken.

- Register the `RolesAllowedDynamicFeature` with the application.
- Make sure you add `Authorizers` when you build your `AuthFilters`.
- Make sure any custom `AuthFilter` you add has the `@Priority(Priorities.AUTHENTICATION)` annotation set (otherwise authorization will be tested before the request's security context is properly set and will fail).
- Annotate the resource *method* with the authorization annotation. Unlike the note earlier in this document that says authorization annotations are allowed on classes, with this poly feature, currently that is not supported. The annotation **MUST** go on the resource *method*

So continuing with the previous example you should add the following configurations

```
... = new BasicCredentialAuthFilter.Builder<BasicPrincipal>()
    .setAuthorizer(new ExampleAuthorizer()).. // set authorizer

... = new OAuthCredentialAuthFilter.Builder<OAuthPrincipal>()
    .setAuthorizer(new ExampleAuthorizer()).. // set authorizer

environment.jersey().register(RolesAllowedDynamicFeature.class);
```

Now we can do

```
@GET
@RolesAllowed({ "ADMIN" })
public Response baseAuthResource(@Auth BasicPrincipal principal) {}

@GET
@RolesAllowed({ "ADMIN" })
public Response oauthResource(@Auth OAuthPrincipal principal) {}
```

Note: The polymorphic auth feature *SHOULD NOT* be used with any other `AuthDynamicFeature`. Doing so may have undesired effects.

DROPWIZARD FORMS

The `dropwizard-forms` module provides you with a support for multi-part forms via Jersey.

11.1 Adding The Bundle

Then, in your application's `initialize` method, add a new `MultiPartBundle` subclass:

```
@Override
public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.addBundle(new MultiPartBundle());
}
```

11.2 Testing

To test resources that utilize multi-part form features, one must add `MultiPartFeature.class` to the `ResourceExtension` as a provider, and register it on the client like the following:

```
@ExtendWith(DropwizardExtensionsSupport.class)
public class MultiPartTest {

    public static final ResourceExtension resourceExtension = ResourceExtension.builder()
        .addProvider(MultiPartFeature.class)
        .addResource(new TestResource())
        .build();

    @Test
    public void testClientMultipart() {
        final FormDataMultiPart multiPart = new FormDataMultiPart()
            .field("test-data", "Hello Multipart");
        final String response = resourceExtension.target("/test")
            .register(MultiPartFeature.class)
            .request()
            .post(Entity.entity(multiPart, multiPart.getMediaType()), String.class);
        assertThat(response).isEqualTo("Hello Multipart");
    }

    @Path("test")
    public static class TestResource {
```

(continues on next page)

(continued from previous page)

```
@POST
@Consumes(MediaType.MULTIPART_FORM_DATA)
public String post(@FormDataParam("test-data") String testData) {
    return testData;
}
}
```

11.3 More Information

For additional and more detailed documentation about the Jersey multi-part support, please refer to the documentation in the [Jersey User Guide](#) and [Javadoc](#).

DROPWIZARD VALIDATION

Dropwizard comes with a host of validation tools out of the box to allow endpoints to return meaningful error messages when constraints are violated. Hibernate Validator is packaged with Dropwizard, so what can be done in Hibernate Validator, can be done with Dropwizard.

12.1 Validations

Almost anything can be validated on resource endpoints. To give a quick example, the following endpoint doesn't allow a null or empty name query parameter.

```
@GET
public String find(@QueryParam("name") @NotEmpty String arg) {
    // ...
}
```

If a client sends an empty or nonexistent name query param, Dropwizard will respond with a 400 Bad Request code with the error: query param name may not be empty.

Additionally, annotations such as HeaderParam, CookieParam, FormParam, etc, can be constrained with violations giving descriptive errors and 400 status codes.

12.1.1 Constraining Entities

If we're accepting client-provided Person, we probably want to ensure that the name field of the object isn't null or blank in the request. We can do this as follows:

```
public class Person {

    @NotEmpty // ensure that name isn't null or blank
    private final String name;

    @JsonCreator
    public Person(@JsonProperty("name") String name) {
        this.name = name;
    }

    @JsonProperty("name")
    public String getName() {
        return name;
    }
}
```

Then, in our resource class, we can add the `@Valid` annotation to the `Person` annotation:

```
@PUT
public Person replace(@NotNull @Valid Person person) {
    // ...
}
```

If the name field is missing, Dropwizard will return a 422 `Unprocessable Entity` response detailing the validation errors: name may not be empty

Note: You don't need `@Valid` when the type you are validating can be validated directly (`int`, `String`, `Integer`). If a class has fields that need validating, then instances of the class must be marked `@Valid`. For more information, see the Hibernate Validator documentation on [Object graphs](#) and [Cascaded validation](#).

Since our entity is also annotated with `@NotNull`, Dropwizard will also guard against null input with a response stating that the body must not be null.

12.1.2 Constraints on optional types

If an entity, field, or parameter is not strictly required, it can be wrapped in an `Optional<T>`, but the wrapped value can still be constrained by setting the annotation parameter `payload = Unwrapping.Unwrap.class`. If the optional container is empty, then the constraints are not applied.

Numeric optional types, such as `OptionalDouble`, `OptionalInt`, and `OptionalLong` do not require explicit unwrapping.

Note: Be careful when using constraints with validation annotations on `Optional<String>` parameters as there is a subtle, but important distinction between null and empty. If a client requests `bar?q=`, the query parameter `q` will evaluate to `Optional.of("")`. If you want `q` to evaluate to `Optional.empty()` in this situation, change the type to `NonEmptyStringParam`.

Note: Parameter types such as `IntParam` and `NonEmptyStringParam` can also be constrained.

There is a caveat regarding `payload = Unwrapping.Unwrap.class` and `*Param` types, as there still are some cumbersome situations when constraints need to be applied to the container and the value.

```
@POST
// The @NotNull is supposed to mean that the parameter is required but the Max(3) is
↪ supposed to
// apply to the contained integer. Currently, this code will fail saying that Max can't
// be applied on an IntParam
public List<Person> createNum(@QueryParam("num")
                             @NotNull(payload = Unwrapping.Unwrap.class)
                             @Max(value = 3, payload = Unwrapping.Unwrap.class)
↪ IntParam num) {
    // ...
}

@GET
// Similarly, the underlying validation framework can't unwrap nested types (an integer
```

(continues on next page)

(continued from previous page)

```

↳wrapped
// in an IntParam wrapped in an Optional), regardless if `Unwrapping.Unwrap.class` is used
public Person retrieve(@QueryParam("num") @Max(3) Optional<IntParam> num) {
    // ...
}

```

To work around these limitations, if the parameter is required check for it in the endpoint and throw an exception, else use `@DefaultValue` or move the `Optional<T>` into the endpoint.

```

@POST
// Workaround to handle required int params and validations
public List<Person> createNum(@QueryParam("num") @Max(3) IntParam num) {
    if (num == null) {
        throw new BadRequestException("query param num must not be null");
    }
    // ...
}

@GET
// Workaround to handle optional int params and validations with DefaultValue
public Person retrieve(@QueryParam("num") @DefaultValue("0") @Max(3) IntParam num) {
    // ...
}

@GET
// Workaround to handle optional int params and validations with Optional
public Person retrieve2(@QueryParam("num") @Max(3) IntParam num) {
    Optional.fromNullable(num);
    // ...
}

```

12.1.3 Enum Constraints

Given the following enum:

```

public enum Choice {
    OptionA,
    OptionB,
    OptionC
}

```

And the endpoint:

```

@GET
public String getEnum(@NotNull @QueryParam("choice") Choice choice) {
    return choice.toString();
}

```

One can expect Dropwizard not only to ensure that the query parameter exists, but to also provide the client a list of valid options query param choice must be one of [OptionA, OptionB, OptionC] when an invalid parameter is provided. The enum that the query parameter is deserialized into is first attempted on the enum's `name()` field and

then `toString()`. During the case insensitive comparisons, the query parameter has whitespace removed with dashes and dots normalized to underscores. This logic is also used when deserializing request body's that contain enums.

12.1.4 Return Value Validations

It's reasonable to want to make guarantees to clients regarding the server response. For example, you may want to assert that no response will ever be null, and if an endpoint creates a `Person` that the person is valid.

```
@POST
@NotNull
@Valid
public Person create() {
    return new Person(null);
}
```

In this instance, instead of returning someone with a null name, Dropwizard will return an HTTP 500 Internal Server Error with the error server response name may not be empty, so the client knows the server failed through no fault of their own.

Analogous to an empty request body, an empty entity annotated with `@NotNull` will return server response may not be null

Warning: Be careful when using return value constraints when endpoints satisfy all of the following:

- Function name starts with `get`
- No arguments
- The return value has validation constraints

If an endpoint satisfies these conditions, whenever a request is processed by the resource that endpoint will be additionally invoked. To give a concrete example:

```
@Path("/")
public class ValidatedResource {
    private AtomicLong counter = new AtomicLong();

    @GET
    @Path("/foo")
    @NotEmpty
    public String getFoo() {
        counter.getAndIncrement();
        return "";
    }

    @GET
    @Path("/bar")
    public String getBar() {
        return "";
    }
}
```

If a `/foo` is requested then `counter` will have increment by 2, and if `/bar` is requested then `counter` will increment by 1. It is our hope that such endpoints are few, far between, and documented thoroughly.

12.2 Limitations

Jersey allows for BeanParam to have setters with *Param annotations. While nice for simple transformations it does obstruct validation, so clients won't receive as instructive of error messages. The following example shows the behavior:

```
@Path("/root")
@Produces(MediaType.APPLICATION_JSON)
public class Resource {

    @GET
    @Path("params")
    public String getBean(@Valid @BeanParam MyBeanParams params) {
        return params.getField();
    }

    public static class MyBeanParams {
        @NotEmpty
        private String field;

        public String getField() {
            return field;
        }

        @QueryParam("foo")
        public void setField(String field) {
            this.field = Strings.nullToEmpty(field).trim();
        }
    }
}
```

A client submitting the query parameter foo as blank will receive the following error message:

```
{"errors":["getBean.arg0.field may not be empty"]}
```

Workarounds include:

- Name BeanParam fields the same as the *Param annotation values
- Supply validation message on annotation: @NotEmpty(message = "query param foo must not be empty")
- Perform transformations and validations on *Param inside endpoint

The same kind of limitation applies for *Configuration* objects:

```
public class MyConfiguration extends Configuration {
    @NotNull
    @JsonProperty("foo")
    private String baz;
}
```

Even though the property's name is foo, the error when property is null will be:

```
* baz may not be null
```

12.3 Annotations

In addition to the annotations defined in Hibernate Validator, Dropwizard contains another set of annotations, which are briefly shown below.

```
public class Person {
    @NotEmpty
    private final String name;

    @NotEmpty
    @OneOf(value = {"m", "f"}, ignoreCase = true, ignoreWhitespace = true)
    // @OneOf forces a value to value within certain values.
    private final String gender;

    @Min(value = 0, payload = Unwrapping.Unwrap.class)
    @Max(value = 10, payload = Unwrapping.Unwrap.class)
    // The integer contained, if present, can attain a min value of 0 and a max of 10.
    private final Optional<Integer> animals;

    @JsonCreator
    public Person(@JsonProperty("name") String name) {
        this.name = name;
    }

    @JsonProperty("name")
    public String getName() {
        return name;
    }

    // Method that must return true for the object to be valid
    @ValidationMethod(message="name may not be Coda")
    @JsonIgnore
    public boolean isNotCoda() {
        return !"Coda".equals(name);
    }
}
```

The reason why Dropwizard defines `@ValidationMethod` is that more complex validations (for example, cross-field comparisons) are often hard to do using declarative annotations. Adding `@ValidationMethod` to any boolean-returning method which begins with `is` is a short and simple workaround:

Note: Due to the rather daft JavaBeans conventions, when using `@ValidationMethod`, the method must begin with `is` (e.g., `#isValidPortRange()`). This is a limitation of Hibernate Validator, not Dropwizard.

12.3.1 Validating Grouped Constraints with @Validated

The @Validated annotation allows for validation groups to be specifically set, instead of the default group. This is useful when different endpoints share the same entity but may have different validation requirements.

Going back to our favorite Person class. Let's say in the initial version of our API, name has to be non-empty, but realized that business requirements changed and a name can't be longer than 5 letters. Instead of switching out the API from unsuspecting clients, we can accept both versions of the API but at different endpoints.

```
// We're going to create a group of validations for each version of our API
public interface Version1Checks { }

// Our second version will extend Hibernate Validator Default class so that any
↪validation
// annotation without an explicit group will also be validated with this version
public interface Version2Checks extends Default { }

public class Person {
    @NotEmpty(groups = Version1Checks.class)
    @Length(max = 5, groups = Version2Checks.class)
    private String name;

    @JsonCreator
    public Person(@JsonProperty("name") String name) {
        this.name = name;
    }

    @JsonProperty
    public String getName() {
        return name;
    }
}

@Path("/person")
@Produces(MediaType.APPLICATION_JSON)
public class PersonResource {

    // For the v1 endpoint, we'll validate with the version1 class, so we'll need to
    ↪change the
    // group of the NotNull annotation from the default of Default.class to
    ↪Version1Checks.class
    @POST
    @Path("/v1")
    public void createPersonV1(
        @NotNull(groups = Version1Checks.class)
        @Valid
        @Validated(Version1Checks.class)
        Person person
    ) {
        // implementation ...
    }

    // For the v2 endpoint, we'll validate with version1 and version2, which implicitly
    // adds in the Default.class.
}
```

(continues on next page)

(continued from previous page)

```
@POST
@Path("/v2")
public void createPersonV2(
    @NotNull
    @Valid
    @Validated({Version1Checks.class, Version2Checks.class})
    Person person
) {
    // implementation ...
}
}
```

Now when clients hit `/person/v1` the `Person` entity will be checked by all the constraints that are a part of the `Version1Checks` group. If `/person/v2` is hit, then all validations are performed.

Warning: If the `Version1Checks` group wasn't set for the `@NotNull` annotation for the `v1` endpoint, the annotation would not have had any effect and a null pointer exception would have occurred when a property of a person is accessed. Dropwizard tries to protect against this class of bug by disallowing multiple `@Validated` annotations on an endpoint that contain different groups.

12.4 Testing

It is critical to test the constraints so that you can ensure the assumptions about the data hold and see what kinds of error messages clients will receive for bad input. The recommended way for testing annotations is through *Testing Resources*, as Dropwizard does a bit of magic behind the scenes when a constraint violation occurs to set the response's status code and ensure that the error messages are user friendly.

```
@Test
public void personNeedsAName() {
    // Tests what happens when a person with a null name is sent to
    // the endpoint.
    final Response post = resources.target("/person/v1").request()
        .post(Entity.json(new Person(null)));

    // Clients will receive a 422 on bad request entity
    assertThat(post.getStatus()).isEqualTo(422);

    // Check to make sure that errors are correct and human readable
    ValidationErrorMessage msg = post.readEntity(ValidationErrorMessage.class);
    assertThat(msg.getErrors())
        .containsOnly("name may not be empty");
}
}
```

12.5 Extending

While Dropwizard provides good defaults for validation error messages, one can customize the response through an `ExceptionHandler<JerseyViolationException>`:

```
/** Return a generic response depending on if it is a client or server error */
public class MyJerseyViolationExceptionHandler implements ExceptionMapper
<JerseyViolationException> {
    @Override
    public Response toResponse(final JerseyViolationException exception) {
        final Set<ConstraintViolation<?>> violations = exception.
getConstraintViolations();
        final Invocable invocable = exception.getInvocable();
        final int status = ConstraintMessage.determineStatus(violations, invocable);
        return Response.status(status)
            .type(MediaType.TEXT_PLAIN_TYPE)
            .entity(status >= 500 ? "Server error" : "Client error")
            .build();
    }
}
```

To register `MyJerseyViolationExceptionHandler` and have it override the default:

```
@Override
public void run(final MyConfiguration conf, final Environment env) {
    env.jersey().register(new MyJerseyViolationExceptionHandler());
    env.jersey().register(new Resource());
}
```

Dropwizard calculates the validation error message through `ConstraintMessage.getMessage`.

If you need to validate entities outside of resource endpoints, the validator can be accessed in the `Environment` when the application is first ran.

```
Validator validator = environment.getValidator();
Set<ConstraintViolation> errors = validator.validate(/* instance of class */)

```


DROPWIZARD VIEWS

The `dropwizard-views-mustache` & `dropwizard-views-freemarker` modules provide you with simple, fast HTML views using either FreeMarker or Mustache.

To enable views for your *Application*, add the `ViewBundle` in the `initialize` method of your `Application` class:

Listing 1: Source file

```
bootstrap.addBundle(new ViewBundle<>());
```

You can pass configuration through to view renderers by overriding `getViewConfiguration`:

Listing 2: Source file

```
bootstrap.addBundle(new ViewBundle<ViewsConfiguration>() {  
    @Override  
    public Map<String, Map<String, String>> getViewConfiguration(ViewsConfiguration  
↪config) {  
        return config.getViewRendererConfiguration();  
    }  
});
```

The returned map should have, for each renderer (such as `freemarker` or `mustache`), a `Map<String, String>` describing how to configure the renderer. Specific keys and their meanings can be found in the FreeMarker and Mustache documentation:

Listing 3: Source file

```
---  
# views: config->mustache  
views:  
  mustache:  
    cache: false  
    # views: config->mustache  
  freemarker:  
    strict_syntax: true
```

Then, in your *resource method*, add a `View` class:

Listing 4: Source file

```
package io.dropwizard.documentation;
```

(continues on next page)

(continued from previous page)

```
import io.dropwizard.views.common.View;

public class PersonView extends View {
    private final Person person;

    public PersonView(Person person) {
        super("person.ftl");
        this.person = person;
    }

    public Person getPerson() {
        return person;
    }
}
```

person.ftl is the path of the template relative to the class name. If this class was com.example.service.PersonView, Dropwizard would then look for the file src/main/resources/com/example/service/person.ftl.

If your template path contains .ftl, .ftlh, or .ftlx, it'll be interpreted as a [FreeMarker](#) template. If it contains .mustache, it'll be interpreted as a Mustache template.

Tip: Dropwizard [Freemarker](#) Views also support localized template files. It picks up the client's locale from their Accept-Language, so you can add a French template in person_fr.ftl or a Canadian template in person_en_CA.ftl.

Your template file might look something like this:

Listing 5: Source file

```
<!-- @ftlvariable name="" type="io.dropwizard.documentation.PersonView" -->
<html lang="en">
<body>
<!-- calls getPerson().getName() and sanitizes it -->
<h1>Hello, ${person.name?html}!</h1>
</body>
</html>
```

The @ftlvariable lets FreeMarker (and any FreeMarker IDE plugins you may be using) know that the root object is a com.example.views.PersonView instance. If you attempt to call a property which doesn't exist on PersonView – getConnectionPool(), for example – it will flag that line in your IDE.

Once you have your view and template, you can simply return an instance of your View subclass:

Listing 6: Source file

```
package io.dropwizard.documentation;

import jakarta.ws.rs.GET;
import jakarta.ws.rs.Path;
import jakarta.ws.rs.PathParam;
import jakarta.ws.rs.Produces;
import jakarta.ws.rs.core.MediaType;
```

(continues on next page)

(continued from previous page)

```

@Path("/people/{id}")
@Produces(MediaType.TEXT_HTML)
public class PersonResource {
    private final PersonDAO dao;

    public PersonResource(PersonDAO dao) {
        this.dao = dao;
    }

    @GET
    public PersonView getPerson(@PathParam("id") String id) {
        return new PersonView(dao.find(id));
    }
}

```

Tip: Jackson can also serialize your views, allowing you to serve both `text/html` and `application/json` with a single representation class.

For more information on how to use FreeMarker, see the [FreeMarker](#) documentation.

For more information on how to use Mustache, see the [Mustache](#) and [Mustache.java](#) documentation.

13.1 Template Errors

By default, if there is an error with the template (eg. the template file is not found or there is a compilation error with the template), the user will receive a `500 Internal Server Error` with a generic HTML message. The exact error will be logged under error mode.

To customize the behavior, create an exception mapper that will override the default one by looking for `ViewRenderException`:

Listing 7: Source file

```

environment.jersey().register(new ExtendedExceptionHandler<WebApplicationException>() {
    @Override
    public Response toResponse(WebApplicationException exception) {
        // Return a response here, for example HTTP 500 (Internal Server Error)
        return Response.serverError().build();
    }

    @Override
    public boolean isMappable(WebApplicationException e) {
        return ExceptionUtils.indexOfThrowable(e, ViewRenderException.class) != -1;
    }
});

```

As an example, to return a 404 instead of an internal server error when one's mustache templates can't be found:

Listing 8: Source file

```
environment.jersey().register(new ExtendedExceptionHandler<WebApplicationException>() {
    @Override
    public Response toResponse(WebApplicationException exception) {
        return Response.status(Response.Status.NOT_FOUND).build();
    }

    @Override
    public boolean isMappable(WebApplicationException e) {
        return ExceptionUtils.getRootCause(e).getClass() == MustacheNotFoundException.
↪class;
    }
});
```

13.2 Caching

By default templates are cached to improve loading time. If you want to disable it during the development mode, set the cache property to false in the view configuration.

Listing 9: Source file

```
views:
  mustache:
    cache: false
```

13.3 Custom Error Pages

To get HTML error pages that fit in with your application, you can use a custom error view. Create a `View` that takes an `ErrorMessage` parameter in its constructor, and hook it up by registering an instance of `ErrorEntityWriter`.

Listing 10: Source file

```
environment.jersey().register(new ErrorEntityWriter<ErrorMessage, View>(MediaType.TEXT_
↪HTML_TYPE, View.class) {
    @Override
    protected View getRepresentation(ErrorMessage errorMessage) {
        return new ErrorView(errorMessage);
    }
});
```

For validation error messages, you'll need to register another `ErrorEntityWriter` that handles `ValidationErrorMessage` objects.

Listing 11: Source file

```
environment.jersey().register(new ErrorEntityWriter<ValidationErrorMessage, View>
↪(MediaType.TEXT_HTML_TYPE, View.class) {
    @Override
```

(continues on next page)

(continued from previous page)

```
protected View getRepresentation(ValidationErrorMessage message) {  
    return new ValidationErrorView(message);  
}  
});
```


DROPWIZARD & SCALA

The `dropwizard-scala` module is now maintained and documented elsewhere.

The `metrics-scala` module is maintained [here](#).

TESTING DROPWIZARD

The `dropwizard-testing` module provides you with some handy classes for testing your representation classes and resource classes. It also provides an extension for JUnit 5.x. A rule for JUnit 4.x is provided by `dropwizard-testing-junit4`

15.1 Testing Representations

While Jackson's JSON support is powerful and fairly easy-to-use, you shouldn't just rely on eyeballing your representation classes to ensure you're producing the API you think you are. You can add unit tests for serializing and deserializing your representation classes to and from JSON.

Let's assume we have a `Person` class which your API uses as both a request entity (e.g., when writing via a PUT request) and a response entity (e.g., when reading via a GET request):

```
public class Person {
    private String name;
    private String email;

    private Person() {
        // Jackson deserialization
    }

    public Person(String name, String email) {
        this.name = name;
        this.email = email;
    }

    @JsonProperty
    public String getName() {
        return name;
    }

    @JsonProperty
    public void setName(String name) {
        this.name = name;
    }

    @JsonProperty
    public String getEmail() {
        return email;
    }
}
```

(continues on next page)

(continued from previous page)

```
@JsonProperty
public void setEmail(String email) {
    this.email = email;
}

// hashCode
// equals
// toString etc.
}
```

15.1.1 Fixtures

First, write out the exact JSON representation of a `Person` in the `src/test/resources/fixtures` directory of your Dropwizard project as `person.json`:

```
{
  "name": "Luther Blissett",
  "email": "lb@example.com"
}
```

15.1.2 Testing Serialization

Next, write a test for serializing a `Person` instance to JSON:

```
import com.fasterxml.jackson.databind.ObjectMapper;
import org.junit.jupiter.api.Test;

import static io.dropwizard.jackson.Jackson.newObjectMapper;
import static org.assertj.core.api.Assertions.assertThat;

class PersonTest {

    private static final ObjectMapper MAPPER = newObjectMapper();

    @Test
    void serializesToJSON() throws Exception {
        final Person person = new Person("Luther Blissett", "lb@example.com");

        final String expected = MAPPER.writeValueAsString(
            MAPPER.readValue(getClass().getResource("/fixtures/person.json"), Person.
↪class));

        assertThat(MAPPER.writeValueAsString(person)).isEqualTo(expected);
    }
}
```

This test uses `AssertJ` assertions and `JUnit` to test that when a `Person` instance is serialized via Jackson it matches the JSON in the fixture file. (The comparison is done on a normalized JSON string representation, so formatting doesn't affect the results.)

15.1.3 Testing Deserialization

Next, write a test for deserializing a `Person` instance from JSON:

```
import com.fasterxml.jackson.databind.ObjectMapper;
import org.junit.jupiter.api.Test;

import static io.dropwizard.jackson.Jackson.newObjectMapper;
import static org.assertj.core.api.Assertions.assertThat;

public class PersonTest {

    private static final ObjectMapper MAPPER = newObjectMapper();

    @Test
    public void deserializesFromJSON() throws Exception {
        final Person person = new Person("Luther Blissett", "lb@example.com");
        assertThat(MAPPER.readValue(getClass().getResource("/fixtures/person.json"),
↳ Person.class))
                .isEqualTo(person);
    }
}
```

This test uses `AssertJ` assertions and `JUnit` to test that when a `Person` instance is deserialized via Jackson from the specified JSON fixture it matches the given object.

15.2 Testing Resources

While many resource classes can be tested just by calling the methods on the class in a test, some resources lend themselves to a more full-stack approach. For these, use `ResourceExtension`, which loads a given resource instance in an in-memory Jersey server:

```
import io.dropwizard.testing.junit5.DropwizardExtensionsSupport;
import io.dropwizard.testing.junit5.ResourceExtension;
import org.junit.jupiter.api.*;
import jakarta.ws.rs.core.Response;
import java.util.Optional;
import static org.assertj.core.api.Assertions.assertThat;
import static org.mockito.Mockito.*;

@ExtendWith(DropwizardExtensionsSupport.class)
class PersonResourceTest {
    private static final PersonDAO DAO = mock(PersonDAO.class);
    private static final ResourceExtension EXT = ResourceExtension.builder()
        .addResource(new PersonResource(DAO))
        .build();
    private Person person;

    @BeforeEach
    void setup() {
        person = new Person();
        person.setId(1L);
    }
}
```

(continues on next page)

(continued from previous page)

```
}

@AfterEach
void tearDown() {
    reset(DAO);
}

@Test
void getPersonSuccess() {
    when(DAO.findById(1L)).thenReturn(Optional.of(person));

    Person found = EXT.target("/people/1").request().get(Person.class);

    assertThat(found.getId()).isEqualTo(person.getId());
    verify(DAO).findById(1L);
}

@Test
void getPersonNotFound() {
    when(DAO.findById(2L)).thenReturn(Optional.empty());
    final Response response = EXT.target("/people/2").request().get();

    assertThat(response.getStatusInfo().getStatusCode()).isEqualTo(Response.Status.
↪NOT_FOUND.getStatusCode());
    verify(DAO).findById(2L);
}
}
```

Instantiate a `ResourceExtension` using its `Builder` and add the various resource instances you want to test via `ResourceExtension.Builder#addResource(Object)`. Use the `@ExtendWith(DropwizardExtensionsSupport.class)` annotation on the class to tell Dropwizard to find any field of type `ResourceExtension`.

In your tests, use `#target(String path)`, which initializes a request to talk to and test your instances.

This doesn't require opening a port, but `ResourceExtension` tests will perform all the serialization, deserialization, and validation that happens inside of the HTTP process.

This also doesn't require a full integration test. In the above *example*, a mocked `PeopleStore` is passed to the `PersonResource` instance to isolate it from the database. Not only does this make the test much faster, but it allows your resource unit tests to test error conditions and edge cases much more easily.

Hint: You can trust `PeopleStore` works because you've got working unit tests for it, right?

15.2.1 Default Exception Mappers

By default, a `ResourceExtension` will register all the default exception mappers (this behavior is new in 1.0). If `registerDefaultExceptionMappers` in the configuration yml is planned to be set to `false`, `ResourceExtension.Builder#setRegisterDefaultExceptionMappers(boolean)` will also need to be set to `false`. Then, all custom exception mappers will need to be registered on the builder, similarly to how they are registered in an `Application` class.

15.2.2 Test Containers

Note that the in-memory Jersey test container does not support all features, such as the `@Context` injection. A different `test container` can be used via `ResourceExtension.Builder#setTestContainerFactory(TestContainerFactory)`.

For example, if you want to use the `Grizzly` HTTP server (which supports `@Context` injections) you need to add the dependency for the Jersey Test Framework providers to your Maven POM and set `GrizzlyWebTestContainerFactory` as `TestContainerFactory` in your test classes.

```
<dependency>
  <groupId>org.glassfish.jersey.test-framework.providers</groupId>
  <artifactId>jersey-test-framework-provider-grizzly2</artifactId>
  <scope>test</scope>
</dependency>
```

```
@ExtendWith(DropwizardExtensionsSupport.class)
class ResourceTestWithGrizzly {
    private static final ResourceExtension EXT = ResourceExtension.builder()
        .setTestContainerFactory(new GrizzlyWebTestContainerFactory())
        .addResource(new ExampleResource())
        .build();

    @Test
    void testResource() {
        assertThat(EXT.target("/example").request()
            .get(String.class))
            .isEqualTo("example");
    }
}
```

15.3 Testing Client Implementations

To avoid circular dependencies in your projects or to speed up test runs, you can test your HTTP client code by writing a JAX-RS resource as test double and let the `DropwizardClientExtension` start and stop a simple Dropwizard application containing your test doubles.

```
@ExtendWith(DropwizardExtensionsSupport.class)
class CustomClientTest {
    @Path("/ping")
    public static class PingResource {
        @GET
        public String ping() {
```

(continues on next page)

(continued from previous page)

```
        return "pong";
    }
}

private static final DropwizardClientExtension EXT = new
↳DropwizardClientExtension(new PingResource());

@Test
void shouldPing() throws IOException {
    final URL url = new URL(EXT.baseUri() + "/ping");
    final String response = new BufferedReader(new InputStreamReader(url.
↳openStream())).readLine();
    assertEquals("pong", response);
}
}
```

Hint: Of course you would use your HTTP client in the @Test method and not `java.net.URL#openStream()`.

The `DropwizardClientExtension` takes care of:

- Creating a simple default configuration.
- Creating a simplistic application.
- Adding a dummy health check to the application to suppress the startup warning.
- Adding your JAX-RS resources (test doubles) to the Dropwizard application.
- Choosing a free random port number (important for running tests in parallel).
- Starting the Dropwizard application containing the test doubles.
- Stopping the Dropwizard application containing the test doubles.

15.4 Integration Testing

It can be useful to start up your entire application and hit it with real HTTP requests during testing. The `dropwizard-testing` module offers helper classes for your easily doing so. The optional `dropwizard-client` module offers more helpers, e.g. a custom `JerseyClientBuilder`, which is aware of your application's environment.

15.4.1 JUnit 5

Adding `DropwizardExtensionsSupport` annotation and `DropwizardAppExtension` extension to your JUnit5 test class will start the app prior to any tests running and stop it again when they've completed (roughly equivalent to having used `@BeforeAll` and `@AfterAll`). `DropwizardAppExtension` also exposes the app's `Configuration`, `Environment` and the app object itself so that these can be queried by the tests.

If you don't want to use the `dropwizard-client` module or find it excessive for testing, you can get access to a Jersey HTTP client by calling the `client` method on the extension. The returned client is managed by the extension and can be reused across tests.

```

@ExtendWith(DropwizardExtensionsSupport.class)
class LoginAcceptanceTest {

    private static DropwizardAppExtension<TestConfiguration> EXT = new
↳DropwizardAppExtension<>(
        MyApp.class,
        ResourceHelpers.resourceFilePath("my-app-config.yaml")
    );

    @Test
    void loginHandlerRedirectsAfterPost() {
        Client client = EXT.client();

        Response response = client.target(
            String.format("http://localhost:%d/login", EXT.getLocalPort()))
            .request()
            .post(Entity.json(loginForm()));

        assertThat(response.getStatus()).isEqualTo(302);
    }
}

```

Warning: Resource classes are used by multiple threads concurrently. In general, we recommend that resources be stateless/immutable, but it's important to keep the context in mind.

15.4.2 Non-JUnit

By creating a `DropwizardTestSupport` instance in your test you can manually start and stop the app in your tests, you do this by calling its `before` and `after` methods. `DropwizardTestSupport` also exposes the app's `Configuration`, `Environment` and the app object itself so that these can be queried by the tests.

```

public class LoginAcceptanceTest {

    public static final DropwizardTestSupport<TestConfiguration> SUPPORT =
        new DropwizardTestSupport<TestConfiguration>(MyApp.class,
            ResourceHelpers.resourceFilePath("my-app-config.yaml"),
            ConfigOverride.config("server.applicationConnectors[0].port", "0") //
↳Optional, if not using a separate testing-specific configuration file, use a randomly
↳selected port
        );

    @BeforeAll
    public void beforeClass() {
        SUPPORT.before();
    }

    @AfterAll
    public void afterClass() {
        SUPPORT.after();
    }
}

```

(continues on next page)

(continued from previous page)

```

@Test
public void loginHandlerRedirectsAfterPost() {
    Client client = new JerseyClientBuilder(SUPPORT.getEnvironment()).build("test_
↪client");

    Response response = client.target(
        String.format("http://localhost:%d/login", SUPPORT.getLocalPort()))
        .request()
        .post(Entity.json(loginForm()));

    assertThat(response.getStatus()).isEqualTo(302);
}
}

```

15.5 Testing Commands

Commands can and should be tested, as it's important to ensure arguments are interpreted correctly, and the output is as expected.

Below is a test for a command that adds the arguments as numbers and outputs the summation to the console. The test ensures that the result printed to the screen is correct by capturing standard out before the command is ran.

```

class CommandTest {
    private final PrintStream originalOut = System.out;
    private final PrintStream originalErr = System.err;
    private final InputStream originalIn = System.in;

    private final ByteArrayOutputStream stdout = new ByteArrayOutputStream();
    private final ByteArrayOutputStream stderr = new ByteArrayOutputStream();
    private Cli cli;

    @BeforeEach
    void setUp() throws Exception {
        // Setup necessary mock
        final JarLocation location = mock(JarLocation.class);
        when(location.getVersion()).thenReturn(Optional.of("1.0.0"));

        // Add commands you want to test
        final Bootstrap<MyConfiguration> bootstrap = new Bootstrap<>(new_
↪MyApplication());
        bootstrap.addCommand(new MyAddCommand());

        // Redirect stdout and stderr to our byte streams
        System.setOut(new PrintStream(stdout));
        System.setErr(new PrintStream(stderr));

        // Build what'll run the command and interpret arguments
        cli = new Cli(location, bootstrap, stdout, stderr);
    }
}

```

(continues on next page)

(continued from previous page)

```

@AfterEach
void teardown() {
    System.setOut(originalOut);
    System.setErr(originalErr);
    System.setIn(originalIn);
}

@Test
void myAddCanAddThreeNumbersCorrectly() {
    final boolean success = cli.run("add", "2", "3", "6");

    SoftAssertions softly = new SoftAssertions();
    softly.assertThat(success).as("Exit success").isTrue();

    // Assert that 2 + 3 + 6 outputs 11
    softly.assertThat(stdOut.toString()).as("stdout").isEqualTo("11");
    softly.assertThat(stdErr.toString()).as("stderr").isEmpty();
    softly.assertAll();
}
}

```

15.6 Testing Database Interactions

In Dropwizard, the database access is managed via the `@UnitOfWork` annotation used on resource methods. In case you want to test database-layer code independently, a `DAOTestExtension` is provided which setups a Hibernate `SessionFactory`.

```

@ExtendWith(DropwizardExtensionsSupport.class)
public class DatabaseTest {

    public DAOTestExtension database = DAOTestExtension.newBuilder()
        .addEntityClass(FooEntity.class).build();

    private FooDAO fooDAO;

    @BeforeEach
    public void setUp() {
        fooDAO = new FooDAO(database.getSessionFactory());
    }

    @Test
    public void createsFoo() {
        FooEntity fooEntity = new FooEntity("bar");
        long id = database.inTransaction(() -> {
            return fooDAO.save(fooEntity);
        });

        assertThat(fooEntity.getId, notNullValue());
    }
}

```

(continues on next page)

(continued from previous page)

```

@Test
public void roundtripsFoo() {
    long id = database.inTransaction(() -> {
        return fooDAO.save(new FooEntity("baz"));
    });

    FooEntity fooEntity = fooDAO.get(id);

    assertThat(fooEntity.getFoo(), equalTo("baz"));
}
}

```

The DAOExtension

- Creates a simple default Hibernate configuration using an H2 in-memory database
- Provides a SessionFactory instance which can be passed to, e.g., a subclass of AbstractDAO
- Provides a function for executing database operations within a transaction

15.7 Testing Configurations

Configuration objects can be tested for correct deserialization and validation. Using the classes created in *polymorphic configurations* as an example, one can assert the expected widget is deserialized based on the type field.

```

public class WidgetFactoryTest {

    private final ObjectMapper objectMapper = Jackson.newObjectMapper();
    private final Validator validator = Validators.newValidator();
    private final YAMLConfigurationFactory<WidgetFactory> factory =
        new YAMLConfigurationFactory<>(WidgetFactory.class, validator, objectMapper,
        ↪ "dw");

    @Test
    public void isDiscoverable() throws Exception {
        // Make sure the types we specified in META-INF gets picked up
        assertThat(new DiscoverableSubtypeResolver().getDiscoveredSubtypes())
            .contains(HammerFactory.class)
            .contains(ChiselFactory.class);
    }

    @Test
    public void testBuildAHammer() throws Exception {
        final WidgetFactory wid = factory.build(new
        ↪ ResourceConfigurationSourceProvider(), "yaml/hammer.yml");
        assertThat(wid).isInstanceOf(HammerFactory.class);
        assertThat(((HammerFactory) wid).createWidget().getWeight()).isEqualTo(10);
    }

    // test for the chisel factory
}

```

If your configuration file contains environment variables or parameters, some additional config is required. As an example, we will use `EnvironmentVariableSubstitutor` on top of a simplified version of the above test.

If we have a configuration similar to the following:

```
widgets:
- type: hammer
  weight: ${HAMMER_WEIGHT:-20}
- type: chisel
  radius: 0.4
```

In order to test this, we would require the following in our test class:

```
public class WidgetFactoryTest {

    private final ObjectMapper objectMapper = Jackson.newObjectMapper();
    private final Validator validator = Validators.newValidator();
    private final YamlConfigurationFactory<WidgetFactory> factory =
        new YamlConfigurationFactory<>(WidgetFactory.class, validator, objectMapper,
        ↪ "dw");

    // test for discoverability

    @Test
    public void testBuildAHammer() throws Exception {
        final WidgetFactory wid = factory.build(new SubstitutingSourceProvider(
            new ResourceConfigurationSourceProvider(),
            new EnvironmentVariableSubstitutor(false)
        ), "yaml/hammer.yaml");
        assertThat(wid).isInstanceOf(HammerFactory.class);
        assertThat(((HammerFactory) wid).createWidget().getWeight()).isEqualTo(20);
    }

    // test for the chisel factory
}
```


UPGRADE NOTES

16.1 Upgrade Notes for Dropwizard 0.7.x

- Update Java source and target versions in `maven-compiler-plugin` to *1.7* (most applications should be already on 1.7);
- Replace Maven dependencies from `com.yammer.dropwizard` to `io.dropwizard`;
- Replace package statements from `com.yammer.dropwizard` to `io.dropwizard` throughout the codebase;
- If you use `dropwizard-db`, update configuration class to use `DataSourceFactory`;
- If you use `dropwizard-hibernate`, update Hibernate bundle by overriding `getDataSourceFactory`;
- If you use `dropwizard-migrations`, update Migrations bundle by overriding `getDataSourceFactory`;
- If you serve static files, add `dropwizard-assets` to dependencies;
- If you use templating, add `dropwizard-views-freemarker` or `dropwizard-views-mustache` accordingly;
- Update the application to override `getName()` instead of providing the bundle with the name;
- Change how resources are added from `environment.addResource(resource)` to `environment.jersey().register(resource)`;
- Once everything is compiling, rename `*Service` class to `*Application`;
- Change test classes extending `ResourceTest` to use `ResourceTestRule`;
- Convert `app.yml` to the new server layout (see `ServerFactory` and `ConnectorFactory`);

16.2 Upgrade Notes for Dropwizard 0.8.x

16.2.1 First

Check out [Migration discussion 0.7.1 to 0.8.0](#) at the `dropwizard-dev` mailing list.

16.2.2 Migration of Apache Commons Lang

The classes were moved to a new package. You have to update the corresponding imports:

search for: `org.apache.commons.lang.`
replace with: `org.apache.commons.lang3.`

16.2.3 Use assertions from AssertJ

Instead of the FEST assertions you should use the AssertJ assertions:

search for: `org.fest.assertions.api.Assertions.`
replace with: `org.assertj.core.api.Assertions.`

16.2.4 Migration of custom URL pattern

If you set a custom URL pattern in your application run method you should move the definition to your configuration file:

Remove from Java code (example):

```
environment.jersey().setUrlPattern("/api/*");
```

Add to configuration file (example):

```
server:  
  rootPath: '/api/*'
```

16.2.5 Migration of Jersey

This is not a simple *search and replace* migration, so I show you a few examples of often used code snippets for integration testing:

Dropwizard Class Rule

The class rule was not modified. It is shown here because it is used in the examples below.

```
@ClassRule  
public static final DropwizardAppRule<SportChefConfiguration> RULE =  
    new DropwizardAppRule<>(App.class, "config.yaml");
```

Executing a GET request

```
final WebTarget target = ClientBuilder.newClient().target(
    String.format("http://localhost:%d/api/user/1", RULE.getLocalPort()));

final Response response = target
    .request(MediaType.APPLICATION_JSON_TYPE)
    .accept(MediaType.APPLICATION_JSON_TYPE)
    .get();

assertThat(response.getStatus()).isEqualTo(Response.Status.OK.getStatusCode());

final User user = response.readEntity(User.class);
assertThat(user.getId()).isEqualTo(1L);
assertThat(user.getFirstName()).isEqualTo("John");
assertThat(user.getLastName()).isEqualTo("Doe");
```

Executing a POST request

```
final WebTarget target = ClientBuilder.newClient().target(
    String.format("http://localhost:%d/api/user", RULE.getLocalPort()));

final User user = new User(0L, "John", "Doe");

final Response response = target
    .request(MediaType.APPLICATION_JSON_TYPE)
    .accept(MediaType.APPLICATION_JSON_TYPE)
    .post(Entity.json(user));

assertThat(response.getStatus()).isEqualTo(Response.Status.CREATED.getStatusCode());

final URI location = response.getLocation();
assertThat(location).isNotNull();

final String path = location.getPath();
final long newId = Long.parseLong(path.substring(path.lastIndexOf("/") + 1));
assertThat(newId).isGreaterThan(0);
```

Executing a empty PUT request

Jersey 2 does not by default allow empty PUT or DELETE requests. If you want to enable this, you have to add a configuration parameter

```
Client client = ClientBuilder.newClient();
client.property(ClientProperties.SUPPRESS_HTTP_COMPLIANCE_VALIDATION, true);
WebTarget target = client.target(
    String.format("http://localhost:%d/api/user", RULE.getLocalPort()));

Response response = target
    .request()
```

(continues on next page)

(continued from previous page)

```
.put(null);  
assertThat(response.getStatus()).isEqualTo(Response.Status.OK.getStatusCode());
```

Request/response filters

If you previously used jersey container filters in your Dropwizard app, `getContainerRequestFilters()` will now fail to resolve:

```
env.jersey()  
.getResourceConfig()  
.getContainerRequestFilters()  
.add(new AuthorizedFilter());
```

You might need to rewrite the filter to JAX-RS 2.0 and then you may use the one and only `.register()` instead.

My filters used imports from `jersey.spi.container` and needed to be rewritten for Jersey 2.x. See also: [Jersey 1.x to 2.x migration guide](#).

```
env.jersey().register(new AuthorizationFilter());
```

16.3 Upgrade Notes for Dropwizard 0.9.x

16.3.1 Migrating Auth

1. Any custom types representing a user need to implement the `Principal` interface
2. In your `Application#run` add

```
environment.jersey().register(RolesAllowedDynamicFeature.class);
```

3. Create an Authorizer

```
public class ExampleAuthorizer implements Authorizer<User> {  
    @Override  
    public boolean authorize(User user, String role) {  
        return user.getName().equals("good-guy") && role.equals("ADMIN");  
    }  
}
```

4. Create an `AuthFilter` using your `Authenticator` and `Authorizer`

```
final BasicCredentialAuthFilter<User> userBasicCredentialAuthFilter =  
    new BasicCredentialAuthFilter.Builder<User>()  
        .setAuthenticator(new ExampleAuthenticator())  
        .setRealm("SUPER SECRET STUFF")  
        .setAuthorizer(new ExampleAuthorizer())  
        .buildAuthFilter();
```

5. Register `AuthDynamicFeature` with your `AuthFilter`

```
environment.jersey().register(new
↳AuthDynamicFeature(userBasicCredentialAuthFilter));
```

6. Register the `AuthValueFactoryProvider.Binder` so with your custom user type if you have one

```
environment.jersey().register(new AuthValueFactoryProvider.Binder(User.class));
```

7. Annotate resources methods that already have `@Auth` with `@RolesAllowed("admin")` where `admin` is a role

```
$ curl 'testUser:secret@localhost:8080/protected'
Hey there, testUser. You know the secret!
```

16.3.2 UnwrapValidatedValue Changes

With the upgrade to Hibernate Validator 5.2.1.Final, the behavior of `@UnwrapValidatedValue` has slightly changed. In some situations, the annotation is now unnecessary. However, when inference is not possible and is ambiguous where the constraint annotation applies, a runtime exception is thrown. This is only a problem when dealing with constraints that can apply to both the wrapper and inner type like `@NotNull`. The fix is to explicitly set `false` or `true` for `@UnwrapValidatedValue`

For instance if you previously had code like:

```
@GET
public String heads(@QueryParam("cheese") @NotNull IntParam secretSauce) {
```

Where `@NotNull` is meant to apply to wrapper type of `IntParam` and not the inner type of `Integer` (as `IntParam` will never yield a null integer). Hibernate Validator doesn't know this, but it does know that `@NotNull` can be applied to both `IntParam` and `Integer`, so in Dropwizard 0.9.x the previous code will now fail and must be changed to

```
@GET
public String heads(@QueryParam("cheese") @NotNull @UnwrapValidatedValue(false) IntParam
↳secretSauce) {
```

For more information on the behavior changes, see [accompanying table for automatic value unwrapping](#)

16.3.3 Logging bootstrap

If you configured console logging in your tests with a utility method shipped with Dropwizard, you should replace calls of `LoggingFactory.bootstrap` to `BootstrapLogging.bootstrap`.

16.4 Upgrade Notes for Dropwizard 1.0.x

16.4.1 Change the project compile and target level to 1.8

Dropwizard 1.0.0 is compiled against JDK 1.8 uses its features extensively. So, to use this version of Dropwizard your project should be compiled and targeted to run on JDK 1.8.

16.4.2 Remove the dropwizard-java8 module

Support for Java 8 features is now provided out of the box.

16.4.3 Migrate dropwizard-spy to dropwizard-http2

If you used the SPDY connector, you should use the HTTP/2 integration now.

```
# - type: spdy3
- type: h2
  port: 8445
  keyStorePath: example.keystore
  keyStorePassword: example
```

16.4.4 Replace Guava's Optional by java.util.Optional in Dropwizard public API

Although Guava's `Optional` should be still supported in your Jersey and JDBI resources, Dropwizard API now exposes optional results as `java.util.Optional`.

For example, in authenticators you should change `Optional.absent` to `Optional.empty`.

16.4.5 Migrate your Hibernate resources to Hibernate 5

Checkout the [Hibernate 5.0 migration guide](#)

16.4.6 Add missing @Valid annotations

In 0.9.x, `@Validated` was sufficient to enable validation. In 1.0.x, it is *necessary* to include `@Valid` as well.

16.5 Upgrade Notes for Dropwizard 1.1.x

Due to [PR #1851](#), users must now add `mockito` as a test dependency

```
<dependency>
  <groupId>org.mockito</groupId>
  <artifactId>mockito-core</artifactId>
  <version>2.7.6</version>
  <scope>test</scope>
</dependency>
```

Else become susceptible to the following error:

```
java.lang.NoClassDefFoundError: org/mockito/Mockito
```

Due to [PR #1695](#), `Cli` no longer allows exceptions to propagate, (which is a net positive), but I did have to rewrite my tests to no longer trap for exceptions but examine `stderr`.

If you used the Hibernate integration, you need to upgrade your data access code to Hibernate 5.2.7 from 5.1.0. Please see the discussion of the change in [PR #1871](#). Please also check out the [Hibernate 5.2 Migration Guide](#).

16.6 Upgrade Notes for Dropwizard 2.0.x

16.6.1 Dropwizard Bill of Materials (BOM)

Starting with Dropwizard 2.0.0, the `io.dropwizard:dropwizard-bom` artifact only specifies the versions of the official Dropwizard modules but no transitive dependencies anymore.

If you want to pin the transitive dependencies, you'll have to use the `io.dropwizard:dropwizard-dependencies` artifact.

It can be used as parent POM, for which you can override individual dependency versions by setting certain Maven properties, see [dropwizard-dependencies/pom.xml](#) for a complete list.

```
<parent>
  <groupId>io.dropwizard</groupId>
  <artifactId>dropwizard-dependencies</artifactId>
  <version>2.0.0</version>
</parent>

<properties>
  <!-- Use older version of Google Guava -->
  <guava.version>28.0-jre</guava.version>
</properties>
```

Alternatively, you can also import it as a regular BOM without the possibility to override specific transitive dependency versions with a Maven property:

```
<dependencyManagement>
  <dependencies>
    <dependency>
      <groupId>io.dropwizard</groupId>
      <artifactId>dropwizard-dependencies</artifactId>
      <version>2.0.0</version>
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

See also: [#2897](#)

16.6.2 Removed Configuration Options

The following configuration options have been removed, so Dropwizard configuration files should no longer use these options

- `soLingerTime`: the configuration option would have become a noop anyways. See [#2490](#) for more info
- `blockingTimeout`: was previously used as an internal jetty failsafe mechanism, [and that use case was no longer deemed necessary](#). If one had previously used `blockingTimeout` to discard slow clients, please use the new configuration options `minRequestDataPerSecond` and `minResponseDataPerSecond`
- `minRequestDataRate`: has been renamed to `minRequestDataPerSecond` and changed from a number to a size like “100 bytes”

16.6.3 Jersey

Dropwizard has upgraded to Eclipse Jersey 2.29, but it has come at some migration cost:

If one created a custom provider (eg: parse / write JSON differently, so a custom `JacksonJaxbJsonProvider` is written), you must annotate the class with the appropriate `@Consumes` and `@Produces` and register it with a Jersey Feature instead of an `AbstractBinder` if it been so previously.

HK2 internal API has been updated, so if you previously had a `AbstractValueFactoryProvider`, that will need to migrate to a `AbstractValueParamProvider`

Jersey Reactive Client API was updated to remove `RxClient`, as rx capabilities are built into the client. You only need to use Dropwizard’s `buildRx` for client when you want a switch the default to something like rxjava 2’s `Flowable`

Context injection on fields in resource instances

The given resource class has different behavior in Dropwizard 1.3 and Dropwizard 2.0 depending on how it is registered.

```
@Path("/")
@Produces(MediaType.APPLICATION_JSON)
public class InfoResource {
    @Context
    UriInfo requestUri;

    @GET
    public String getInfo() {
        return requestUri.getRequestUri().toString()
    }
}
```

There are two ways to register this resource:

```
@Override
public void run(InfoConfiguration configuration, Environment environment) {
    // 1. Register an instance of the resource
    environment.jersey().register(new InfoResource());

    // 2. Register the class as a resource
    environment.jersey().register(InfoResource.class);
}
```


The first method (registering an instance) will now not work in Dropwizard 2.0. Migrating resource instances with field context injections to Dropwizard 2.0 involves pushing the field into a parameter in the desired endpoint

```

@Path("/")
@Produces(MediaType.APPLICATION_JSON)
public class InfoResource {
-   @Context
-   UriInfo requestUri;

    @GET
-   public String getInfo() {
+   public String getInfo(@Context UriInfo requestUri) {
        return requestUri.getRequestUri().toString()
    }
}

```

For more information see [#2781](#)

16.6.4 More Secure TLS

Dropwizard 2.0, by default, only allows cipher suites that support forward secrecy. The only cipher suites newly disabled are those under the `TLS_RSA_*` family. Clients who don't support forward secrecy (expected to be a small amount) may now find that they can't communicate with a Dropwizard 2.0 server. If necessary one can override what cipher suites are blacklisted using the `excludedCipherSuites` configuration option.

Dropwizard 2.0, by default, only supports TLS 1.2. While Dropwizard 1.x effectively only supported TLS 1.2, due to the supported cipher suites, one could still conceivably configure their server or receive a client that could negotiate a TLS 1.0 or 1.1 connection. One can still decide what TLS protocols are on the blacklist by configuring `excludedProtocols`

We also hope that in 2.0 it is more clear what protocols and cipher suites are enabled / disabled, as previously one would see the following statement logged on startup:

```
Supported protocols: [SSLv2Hello, SSLv3, TLSv1, TLSv1.1, TLSv1.2]
```

While not technically wrong, displaying the protocols that *could* be enabled is misleading as it makes one believe that Dropwizard employs extremely unsafe defaults. We've reworked what is logged to only the protocols and cipher suites that Dropwizard *will* expose. And log the protocols and cipher suites that Dropwizard will reject, and thus could expose them if configured to do so. So now you'll see the following in the logs:

```
Enabled protocols: [TLSv1.2]
Disabled protocols: [SSLv2Hello, SSLv3, TLSv1, TLSv1.1]
```

16.6.5 Jackson Changes

`DeserializationFeature.FAIL_ON_UNKNOWN_PROPERTIES` is now disabled by default, so unrecognized fields will now be silently ignored. One can revert back to the 1.x behavior with:

```

public void initialize(Bootstrap<ExampleConfiguration> bootstrap) {
    bootstrap.getObjectMapper().enable(DeserializationFeature.FAIL_ON_UNKNOWN_
↪PROPERTIES);
}

```

16.6.6 Support for JDBI 2.x moved out of Dropwizard core modules

The `dropwizard-jdbi` module has been moved out of Dropwizard core modules (#2922).

The reason for this is that JDBI 2.x hasn't been updated since January 2017 and the `dropwizard-jdbi3` module, which targets its successor Jdbi 3.x, still is part of the Dropwizard core modules.

If you want to keep using JDBI 2.x, you can change the Maven coordinates of `dropwizard-jdbi` as follows:

```
<!-- Old artifact coordinates -->
<dependency>
  <groupId>io.dropwizard</groupId>
  <artifactId>dropwizard-jdbi</artifactId>
  <version>2.0.0</version>
</dependency>
```

```
<!-- New artifact coordinates -->
<dependency>
  <groupId>io.dropwizard.modules</groupId>
  <artifactId>dropwizard-jdbi</artifactId>
  <version>2.0.0</version>
</dependency>
```

16.6.7 Miscellaneous

Improved validation message for min/max duration

`@MinDuration` / `@MaxDuration` have had their validation messages improved, so instead of

messageRate must be less than (or equal to, if in 'inclusive' mode) 1 MINUTES

one will see if inclusive is true

```
messageRate must be less than or equal to 1 MINUTES
```

if inclusive is false:

```
messageRate must be less than 1 MINUTES
```

Task execute method

The `parameters` argument of the `Task.execute` method has a slightly different `Map` type. Classes extending the abstract class `Task` should therefore change

```
@Override
public void execute(ImmutableMultimap<String, String> parameters, PrintWriter output)
    throws Exception {
```

into

```
@Override
public void execute(Map<String, List<String>> parameters, PrintWriter output) throws
    Exception {
```

16.7 Upgrade Notes for Dropwizard 2.1.x

16.7.1 Hibernate Validator Expression Language overhaul

The Expression Language is now disabled by default for custom violations.

The `@SelfValidating` feature has been a source of various security vulnerabilities in Dropwizard:

- [CVE-2020-5245](#)
- [CVE-2020-11002](#)

This originally led to the introduction of `SelfValidating#escapeExpressions()`.

Due to some changes regarding the interpolation of messages in custom violations in Hibernate Validator 6.2.0.Final and later, this flag is not required anymore and has been removed in Dropwizard 2.1.0.

While strongly discouraged, you can enable EL interpolation in custom violations with Hibernate Validator, by customizing the `HibernateValidatorConfiguration` created by `Validators#newConfiguration()` and use the Hibernate Validator instance created from it in your application via `Environment#setValidator(Validator)`.

Details about these changes in Hibernate Validator can be found at:

- [Hibernate Validator 6.2.0.Final - 12.9. Enabling Expression Language features](#)
- [HV-1816: Disable Expression Language by default for custom constraint violations](#)
- [hibernate/hibernate-validator#1138](#)

16.7.2 Migrating from dropwizard-health

Migrating from the HealthBundle

`dropwizard/dropwizard-health` module Starting in Dropwizard 2.1, the functionality from the external `dropwizard/dropwizard-health` module has been integrated into the main `dropwizard/dropwizard` project.

Perform the following steps to migrate your application:

- Remove the `HealthBundle` from your Application's `initialize()` method.
- Remove the reference to `HealthConfiguration` from your application's `Configuration` class.
- Migrate YAML configuration previously associated with your `HealthConfiguration` field to instead be nested under a top-level `health` field (at the same level as `server`, `logging`, and `metrics`).
- Review the new and changed configuration options mentioned below.
- Ensure that you don't already have a custom configuration class member using the `health` name or rename it to something else or nest it under another field in order not to conflict with the new top-level field.
- Remove the dependency on the `io.dropwizard.modules:dropwizard-health` library from your application's `pom.xml` file (or other build dependency definition file).

New Configuration Options

- `enabled` flag added to turn on/off health functionality (defaults to `true`).
- `name` configuration added, used in metric naming, thread naming, and log messages.
- `responder` factory added, to control how health check requests are responded to. For more info, see *the config reference*.
- `responseProvider` factory added, to control the body of health check responses, separate from the mechanics of responding. For more info, see *the config reference*.

Changed Configuration Options

- `delayedShutdownHandlerEnabled` default value changed to `false`.
- `servlet` removed in favor of the new `responder` and `responseProvider` factories.

16.7.3 Changes in versioning

Although Dropwizard tries to postpone big changes to major releases, some breaking changes had to be introduced into Dropwizard 2.1. This change is necessary due to new versioning of Dropwizard releases. The Dropwizard 2.x releases will stay on a Java 8 baseline and the `javax` namespace. Dropwizard 3.x will stay on the `javax` namespace too, but will drop support for Java 8 and upgrade to Java 11 instead. Dropwizard 4.x will eventually mirror the 3.x versions on the `jakarta` namespace and maybe introduce some more changes.

Therefore major updates for the Java 8 baseline have to be brought to minor releases on the 2.x branch. The major changes introduced in 2.1 are the following:

Library	Version change
argparse	0.8.x → 0.9.x
Hibernate Validator	6.1.x → 6.2.x
Jackson	2.10.x → 2.13.x
Jersey	2.33 → 2.35
Liquibase	3.10.x → 4.9.x
Dropwizard Metrics	4.1.x → 4.2.x

Upgrade to Liquibase 4.x

Most of the updates come with low migration cost, but Liquibase gets a major version upgrade and needs some attention.

Liquibase 4.x changes the way it finds files. This means previously recognized migration files could be reported as missing. Liquibase lets users of the library specify paths, where it should search for files. Dropwizard therefore adds the file system specific roots to these `root` paths, as well as the code location (of the current JAR).

This essentially means migration files now must be specified with absolute paths or be located under `src/main/resources` and specified relative to that path.

Upgrade to Jersey 2.35

The upgrade of Jersey from version 2.33 to 2.35 introduces a behavior change in the handling of `Optional<T>` parameters. If such a parameter is invalid, now a status code `404` is returned instead of the former `400` status code.

16.7.4 Jackson Blackbird as default

Dropwizard is now registering the [Jackson Blackbird](#) module. This is the recommended setup for Java 9 and later.

If the [Jackson Afterburner](#) module is on the class path, it will be preferred over the [Jackson Blackbird](#) module.

```
<dependency>
  <groupId>com.fasterxml.jackson.module</groupId>
  <artifactId>jackson-module-afterburner</artifactId>
  <!-- Unnecessary when using dropwizard-dependencies -->
  <version>${jackson.version}</version>
</dependency>
```

16.7.5 Modification of the client in DropwizardAppExtension

The `DropwizardAppExtension` previously obtained a Jersey client configured with the default `HttpURLConnectionProvider`. To support HTTP methods, which are not implemented in the `HttpURLConnection` class, workarounds are enabled. These workarounds modify the JDK classes to support the requested method. Starting from Java 16, it is no longer possible to modify JDK classes. Therefore workarounds can cause problems on Java 16+.

Dropwizard 2.1.1 changes the Jersey client returned by the `DropwizardAppExtension#client()` method. Now the `GrizzlyConnectorProvider` is registered by default to use the Grizzly Async Http Client instead of the `HttpURLConnection`. Using the Grizzly client should cause no problems when migrating to Dropwizard 2.1.1.

However, if you are experiencing issues with the client modifications, you can easily instantiate a client without the new connector:

```
new JerseyClientBuilder()
  .register(new JacksonFeature(getObjectMapper()))
  .property(ClientProperties.CONNECT_TIMEOUT, DEFAULT_CONNECT_TIMEOUT_MS)
  .property(ClientProperties.READ_TIMEOUT, DEFAULT_READ_TIMEOUT_MS)
  .property(URLConnectionProvider.SET_METHOD_WORKAROUND, true)
  .build();
```

The property `URLConnectionProvider.SET_METHOD_WORKAROUND` is optional and should be avoided on Java versions after Java 15.

16.7.6 Allow parameter conversion exceptions to be mapped to responses by exception mappers

Some parameter validation logic has been fixed to allow for custom remapping of error messages. As a result, the API responses may be a little different than before in certain scenarios.

The responses may be different than before this change if the default exception mappers are disabled, or custom exception mappers are registered:

- If the default exception mappers provided by Dropwizard are disabled (`registerDefaultExceptionMappers` option is off) and no other exception mappers are configured, then the response will have the same status code as before this change, but with no entity or media type.

- If an exception mapper is registered which takes precedence over `LoggingExceptionHandler` for mapping `WebApplicationException` instances, then the response will be generated by that exception mapper. The `WebApplicationException` instance received by the exception mapper will have the same status code and message as the response had before this change.

In addition, the `AbstractParam.mediaType()` method has been deprecated, and the responses may not use the media type returned by that method.

There are two types of parameters that may result in different error responses in the above situations if there is a validation error:

- Enum parameters
- Parameters that are any subclass of `AbstractParam` (Dropwizard provides the subclasses `InstantParam`, `InstantSecondParam`, `LocalDateParam`, `LocalDateTimeParam`, `LocalTimeParam`, `OffsetDateTimeParam`, `YearMonthParam`, `YearParam`, `ZonedDateTimeParam`, `ZoneIdParam`, `BooleanParam`, `DateTimeParam`, `DurationParam`, `InstantParam`, `IntParam`, `LocalDateParam`, `LongParam`, `NonEmptyStringParam`, `SizeParam`, and `UUIDParam`)

If this code change causes an undesired change in the error responses:

- If the error responses are not being mapped to a response by any exception mapper, then either register the default exception mappers (by enabling the `registerDefaultExceptionMappers` option) or register Dropwizard's `LoggingExceptionHandler` to restore the previous behavior, or register your own custom exception mapper that maps `WebApplicationException` instances.
- If you have a custom subclass of `AbstractParam` which overrides `AbstractParam.mediaType()`, then the media type should be set in the exception mapper which maps `WebApplicationException` instances instead. It is also possible to override `AbstractParam.error(String input, Exception e)` to return a response with an entity and the correct media type directly.

See [PR #5691](#) for more details about this change.

16.8 Upgrade Notes for Dropwizard 3.0.x

16.8.1 Java version changes

As already mentioned in the [Upgrade notes for 2.1.x](#), Dropwizard 3.0.0 drops support for Java versions less than 11.

16.8.2 Jetty 10

The main change introduced in Dropwizard 3.0.0 is the upgrade to Jetty 10.0.x. Jetty 10.0.x is built for Java 11, therefore the Java version change was necessary.

This change comes with some migration cost. For detailed information regarding the changes introduced in Jetty 10.0.x, refer to the [Jetty migration guide](#).

The main changes for Dropwizard are the following:

- the `addLifecycleListener(...)` method was replaced by the `addEventListener(...)` method
- the `SecureRequestCustomizer` was changed (see [SNI host checking](#))
- the option to exclude specific user agents got removed from the `GzipHandlerFactory`

SNI host checking

Jetty 10.0.x introduces stricter SNI host checking. Therefore you may encounter problems when making requests over HTTPS.

To solve this issue, the `HttpsConnectorFactory` got the `disableSniHostCheck` configuration option, which defaults to `false` to enable strict security for an application. When setting it to `true`, the SNI host check gets disabled.

16.8.3 Apache HttpClient 5

The Apache HttpClient was updated to version 5.x. This version moves several classes to other packages, moves classes between HttpClient and HttpCore and changes all namespace names. The new prefixes are `org.apache.http.client5` for the client classes and `org.apache.http.core5` for the core classes.

The most functions from Dropwizard are provided as before, but some changes have to be made:

- the `normalizeUri` setting is removed
- the `CredentialsProvider` is replaced by the `CredentialsStore`
- the `ServiceUnavailableRetryStrategy` is removed
- the `HttpRequestRetryHandler` is replaced by the `HttpRequestRetryStrategy`

For more information refer to the [Apache HttpClient 5.0 migration guide](#).

16.8.4 Dropwizard Package Structure and JPMS

In order to properly support the Java Platform Module System (JPMS), the Java packages in modules must not overlap, or put differently, the packages may not be split into multiple modules.

Dropwizard 3.0.0 won't enable full support for the JPMS. Instead, as a transition step, automatic modules are introduced.

See also: [Java 9 Migration Guide: Split Packages](#)

Starting with Dropwizard 3.0.0 some core classes have been moved into distinct and clearly split up packages to clean up overlapping package structure.

Affected packages:

Maven module	Old package	New package
dropwizard-core	<code>io.dropwizard</code>	<code>io.dropwizard.core</code>
dropwizard-logging	<code>io.dropwizard.logging</code>	<code>io.dropwizard.logging.common</code>
dropwizard-metrics	<code>io.dropwizard.metrics</code>	<code>io.dropwizard.metrics.common</code>
dropwizard-views	<code>io.dropwizard.views</code>	<code>io.dropwizard.views.common</code>

This means that imports of core classes such as `io.dropwizard.Application` have to be updated for the new package structure and refer to `io.dropwizard.core.Application`.

16.8.5 JUnit 4.x support

Support for testing with JUnit 4.x has been moved from *dropwizard-testing* to *dropwizard-testing-junit4*.

16.8.6 Removal of classes from *dropwizard-util*

Many classes from the *dropwizard-util* module are now obsolete, since the Java standard library provides replacements for them.

For example the `Sets` class provided helper methods for creating `Set` instances. Starting from Java 9, this can be done by using `Set.of(...)`.

16.8.7 Jadira Usertype Core library

Dropwizard previously registered the types from the Jadira Usertype Core library automatically. In order to align the versions 3.0.x and 4.0.x we will drop support for this library in Dropwizard 3 because the current version of the Jadira Usertype Core library doesn't support Hibernate 6.x, which will be used in Dropwizard 4.0.x.

If you want to continue using this library, you have to set the property `jadira.usertype.autoRegisterUserTypes` to `true` in your application's database configuration and add a dependency on the current version of the Usertype Core library.

16.8.8 Logback layout conversion words

Logback allows to use specific *conversion words* in its `PatternLayout` to insert information obtained by an instance of a specific `Converter`.

Previously, Dropwizard has overridden the abbreviated conversion words for exceptions (`ex`, `xEx` and `rEx`) to apply stack trace prefixing with a `!` rather than a `tab`.

In Dropwizard 3.x these overrides are removed and all exception conversion words work as documented in the Logback manual. To apply the stack trace prefixing, new conversion words are introduced with a prefix `dw`. Therefore the following new conversion words can be used:

- `dwEx`, `dwException` and `dwThrowable` instead of `ex`, `exception` and `throwable`
- `dwXEx`, `dwXException` and `dwXThrowable` instead of `xEx`, `xException` and `xThrowable`
- `dwREx` and `dwRootException` instead of `rEx` and `rootException`

Those newly introduced conversion words work like the Logback ones, except that the first `tab` is replaced by a `!`.

To simplify the upgrade to Dropwizard 3.x for most users, the default Dropwizard logging layout is modified to use the new Dropwizard specific conversion words.

16.9 Upgrade Notes for Dropwizard 4.0.x

16.9.1 Transition to Jakarta EE

For now, all previously released Dropwizard versions used Java/Jakarta EE dependencies under the `javax` namespace. Dropwizard 4.0.x will now transition to Jakarta EE 9 components and therefore utilize the new `jakarta` namespace for many components.

That basically means that most of the imports of existing applications using Dropwizard 3.0.x will have to be changed from the `javax` to the `jakarta` namespace. However, other components still use the `javax` namespace, so a simple search and replace could break other imports.

Jakarta EE compatibility

As stated above, Dropwizard 4.0.x will transition to Jakarta EE 9 components. This means, Dropwizard will now try to be consistent with one specific EE version. Therefore Dropwizard 4.0.x will stay on components of Jakarta EE 9 and the transition to components of Jakarta EE 10 will be postponed to an other release series. Since the EE version bump will probably introduce breaking changes, the Jakarta EE 10 components will be most likely integrated in Dropwizard 5.0.x and not in a 4.1.x release.

16.9.2 Features of Dropwizard 4.0.x

Dropwizard 4.0.x will include all features of Dropwizard 3.0.x as we plan to keep the versions (mostly) in sync. But since the Jakarta EE library development now continues on the `jakarta` namespace, new features may be available only for Dropwizard 4.0.x. Dropwizard 4.0.x may therefore include additional features that are not included in Dropwizard 3.0.x. One such example is the upgrade to Hibernate 6.

16.9.3 Hibernate 6

Hibernate 5.6 provides compatible implementations for JPA 2.2 and for Jakarta Persistence 3.0. But Hibernate 6.0 and 6.1 still provide compatible implementations for Jakarta Persistence 3.0, so we upgraded to Hibernate 6.1 in Dropwizard 4.0.x.

This introduces the following changes:

- removal of `Criteria`: all methods taking `Criteria` instances as parameters are removed since it isn't supported by Hibernate 6 any more
- removal of the restriction for `Serializable` keys: Hibernate reworked its type system and now every `Object` can be a key. Therefore parameters are changed from `Serializable` to `Object`
- removal of `AvailableSettings.USE_NEW_ID_GENERATOR_MAPPINGS`: Dropwizard already used the default value and now this property isn't set any more

If any other aspects from Hibernate are used in an application rather than those provided by Dropwizard, there might be an additional migration cost. Please follow the [Hibernate 6 migration guide](#), if you encounter any problems.

Hibernate 6.0 has already reached its end-of-life, so Dropwizard 4.0.x includes Hibernate 6.1. You may need to follow the [Hibernate 6.1 migration guide](#) as well.

DROPWIZARD EXAMPLE, STEP BY STEP

The `dropwizard-example` module provides you with a working Dropwizard Example Application.

- Preconditions
 - Make sure you have [Maven](#) installed
 - Make sure `JAVA_HOME` points at JDK 11
 - Make sure you have `curl`
- Preparations to start the Dropwizard Example Application
 - Open a terminal / cmd
 - Navigate to the project folder of the Dropwizard Example Application
 - `mvn clean install`
 - `java -jar target/dropwizard-example-1.0.0.jar db migrate example.yml`
 - The statement above ran the liquibase migration in `/src/main/resources/migrations.xml`, creating the table schema
- Starting the Dropwizard Example Application
 - You can now start the Dropwizard Example Application by running `java -jar target/dropwizard-example-1.0.0.jar server example.yml`
 - Alternatively, you can run the Dropwizard Example Application in your IDE: `com.example.helloworld.HelloWorldApplication server example.yml`
- Working with the Dropwizard Example Application
 - Insert a new person: `curl -H "Content-Type: application/json" -d '{"fullName":"John Doe", "jobTitle" : "Chief Wizard" }' http://localhost:8080/people`
 - Retrieve that person: `curl http://localhost:8080/people/1`
 - View that person in a freemarker template: curl or open in a browser `http://localhost:8080/people/1/view_freemarker`
 - View that person in a mustache template: curl or open in a browser `http://localhost:8080/people/1/view_mustache`

DROPWIZARD CONFIGURATION REFERENCE

18.1 Servers

Tweaking some of the options will require good understanding of how Jetty is working. See the [Jetty architecture chapter](#) for reference.

```
server:  
  type: default  
  maxThreads: 1024
```


18.1.1 All

Name	Default	Description
type	default	<ul style="list-style-type: none"> • default • simple
maxThreads	1024	The maximum number of threads the thread pool is allowed to grow. Jetty will throw <code>java.lang.IllegalStateException: Insufficient threads:</code> in case of too aggressive limit on the thread count.
minThreads	8	The minimum number of threads to keep alive in the thread pool. Note that each Jetty connector consumes threads from the pool. See <i>HTTP connector</i> how the thread counts are calculated.
maxQueuedRequests	1024	The maximum number of requests to queue before blocking the acceptors.
responseMeteredLevel	COARSE	The response metered level to decide what response code meters are included
metricPrefix	(none)	The metricPrefix to use in the metric name for jetty metrics
idleThreadTimeout	1 minute	The amount of time a worker thread can be idle before being stopped.
nofileSoftLimit	(none)	The number of open file descriptors before a soft error is issued. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
nofileHardLimit	(none)	The number of open file descriptors before a hard error is issued. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
gid	(none)	The group ID to switch to once the connectors have started. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
uid	(none)	The user ID to switch to once the connectors have started. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
user	(none)	The username to switch to once the connectors have started. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
group	(none)	The group to switch to once the connectors have started. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
umask	(none)	The umask to switch to once the connectors have started. Requires
18.1. Servers		Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .
startsAsRoot	(none)	Whether or not the Dropwizard application is started as a root user. Requires Jetty's <code>libsetuid.so</code> on <code>java.library.path</code> .

GZip

```
server:
  gzip:
    bufferSize: 8KiB
```

Name	De- fault	Description
enabled	true	If true, all requests with <code>gzip</code> in the <code>Accept-Encoding</code> header will have their response entities compressed and requests with <code>gzip</code> in the <code>Content-Encoding</code> header will have their request entities decompressed.
minimumEntity- Size	256 bytes	All response entities under this size are not compressed.
bufferSize	8KiB	The size of the buffer to use when compressing.
exclude- dUserA- gentPat- terns	[]	The set of user agent patterns to exclude from compression.
com- pressed- Mime- Types	Jetty's de- fault	The list of mime types to compress. The default is all types apart the commonly known image, video, audio and compressed types.
included- Methods	Jetty's de- fault	The list of HTTP methods to compress. The default is to compress only GET responses.
deflate- Compres- sionLevel	-1	The compression level used for deflation(compression).
gzipCom- patibleIn- flation	true	This option is unused and deprecated as compressed requests without header info are unsupported
syncFlush	false	The flush mode. Set to true if the application wishes to stream (e.g. SSE) the data, but this may hurt compression performance (as all pending output is flushed).

Request Log

The new request log uses the `logback-access` library for processing request logs, which allow to use an extended set of logging patterns. See the `logback-access-pattern` docs for the reference.

```
server:
  requestLog:
    appenders:
      - type: console
```

Name	Default	Description
appen- ders	console ap- pender	The set of AppenderFactory appenders to which requests will be logged. See <i>logging</i> for more info.

Classic Request Log

The classic request log uses the `logback-classic` library for processing request logs. It produces logs only in the standard NCSA common log format, but allows to use an extended set of appenders.

```
server:
  requestLog:
    type: classic
    timeZone: UTC
    appenders:
      - type: console
```

Name	Default	Description
time-Zone	UTC	The time zone to which request timestamps will be converted.
appenders	console appender	The set of AppenderFactory appenders to which requests will be logged. See <i>logging</i> for more info.

Server Push

Server push technology allows a server to send additional resources to a client along with the requested resource. It works only for HTTP/2 connections.

```
server:
  serverPush:
    enabled: true
    associatePeriod: '4 seconds'
    maxAssociations: 16
    refererHosts: ['dropwizard.io', 'dropwizard.github.io']
    refererPorts: [8444, 8445]
```

Name	Default	Description
enabled	false	If true, the filter will organize resources as primary resources (those referenced by the <code>Referer</code> header) and secondary resources (those that have the <code>Referer</code> header). Secondary resources that have been requested within a time window from the request of the primary resource will be associated with it. The next time a client requests the primary resource, the server will send to the client the secondary resources along with the primary in a single response.
associatePeriod	4 seconds	The time window within which a request for a secondary resource will be associated to a primary resource.
maxAssociations	16	The maximum number of secondary resources that may be associated to a primary resource.
referrerHosts	All hosts	The list of referrer hosts for which the server push technology is supported.
referrerPorts	All ports	The list of referrer ports for which the server push technology is supported.

18.1.2 Simple

Extends the attributes that are available to *all servers*

```

server:
  type: simple
  applicationContextPath: /application
  adminContextPath: /admin
  connector:
    type: http
    port: 8080

```

Name	Default	Description
connector	http connector	HttpConnectorFactory HTTP connector listening on port 8080. The ConnectorFactory connector which will handle both application and admin requests. TODO link to connector below.
applicationContextPath	/application	The context path of the application servlets, including Jersey.
adminContextPath	/admin	The context path of the admin servlets, including metrics and tasks.

18.1.3 Default

Extends the attributes that are available to *all servers*

```
server:
  adminMinThreads: 1
  adminMaxThreads: 64
  adminContextPath: /
  applicationContextPath: /
  applicationConnectors:
    - type: http
      port: 8080
    - type: https
      port: 8443
      keyStorePath: example.keystore
      keyStorePassword: example
      validateCerts: false
  adminConnectors:
    - type: http
      port: 8081
    - type: https
      port: 8444
      keyStorePath: example.keystore
      keyStorePassword: example
      validateCerts: false
```

Name	Default	Description
application-Connectors	An HTTP connector listening on port 8080.	A set of <i>connectors</i> which will handle application requests.
adminConnectors	An HTTP connector listening on port 8081.	An HTTP connector listening on port 8081. A set of <i>connectors</i> which will handle admin requests.
admin-MinThreads	1	The minimum number of threads to use for admin requests.
adminMax-Threads	64	The maximum number of threads to use for admin requests.
adminContextPath	/	The context path of the admin servlets, including metrics and tasks.
application-ContextPath	/	The context path of the application servlets, including Jersey.

18.2 Connectors

18.2.1 HTTP

```
# Extending from the default server configuration
server:
  applicationConnectors:
    - type: http
      port: 8080
      bindHost: 127.0.0.1 # only bind to loopback
```

(continues on next page)

(continued from previous page)

```
inheritChannel: false
headerCacheSize: 512 bytes
outputBufferSize: 32KiB
maxRequestHeaderSize: 8KiB
maxResponseHeaderSize: 8KiB
inputBufferSize: 8KiB
idleTimeout: 30 seconds
minBufferPoolSize: 64 bytes
bufferPoolIncrement: 1KiB
maxBufferPoolSize: 64KiB
minRequestDataPerSecond: '0 bytes'
minResponseDataPerSecond: '0 bytes'
acceptorThreads: 1
selectorThreads: 2
acceptQueueSize: 1024
reuseAddress: true
useServerHeader: false
useDateHeader: true
useForwardedHeaders: false
useProxyProtocol: false
httpCompliance: RFC7230
```

Name	Default	Description
port	8080	The TCP/IP port on which to listen for incoming connections.
bindHost	(none)	The hostname to bind to.
inheritChannel	false	Whether this connector uses a channel inherited from the JVM. Use it with <code>Server::Starter</code> , to launch an instance of Jetty on demand.
headerCacheSize	512 bytes	The size of the header field cache.
outputBufferSize	32KiB	The size of the buffer into which response content is aggregated before being sent to the client. A larger buffer can improve performance by allowing a content producer to run without blocking, however larger buffers consume more memory and may induce some latency before a client starts processing the content.
maxRequestHeaderSize	8KiB	The maximum allowed size in bytes for the HTTP request line and HTTP request headers. Larger headers will allow for more and/or larger cookies plus larger form content encoded in a URL. However, larger headers consume more memory and can make a server more vulnerable to denial of service attacks.
maxResponseHeaderSize	8KiB	The maximum size of a response header. Larger headers will allow for more and/or larger cookies and longer HTTP headers (eg for redirection). However, larger headers will also consume more memory.
inputBufferSize	8KiB	The size of the per-connection input buffer.
idleTimeout	30 seconds	The maximum idle time for a connection, which roughly translates to the <code>java.net.Socket#setSoTimeout(int)</code> call, although with NIO implementations other mechanisms may be used to implement the timeout. The max idle time is applied when waiting for a new message to be received on a connection or when waiting for a new message to be sent on a connection. This value is interpreted as the maximum time between some progress being made on the connection. So if a single byte is read or written, then the timeout is reset.
minBufferPoolSize	64 bytes	The minimum size of the buffer pool.
bufferPoolIncrement	1KiB	The increment by which the buffer pool should be increased.
maxBufferPoolSize	64KiB	The maximum size of the buffer pool.

18.2.2 HTTPS

Extends the attributes that are available to the *HTTP connector*

```
# Extending from the default server configuration
server:
  applicationConnectors:
    - type: https
      port: 8443
      ....
      keyStorePath: /path/to/file
      keyStorePassword: changeit
      keyStoreType: JKS
      keyStoreProvider:
      trustStorePath: /path/to/file
      trustStorePassword: changeit
      trustStoreType: JKS
      trustStoreProvider:
      keyManagerPassword: changeit
      needClientAuth: false
      wantClientAuth:
      certAlias: <alias>
      crlPath: /path/to/file
      enableCRLDP: false
      enableOCSP: false
      maxCertPathLength: (unlimited)
      ocsponderUrl: (none)
      jceProvider: (none)
      validateCerts: false
      validatePeers: false
      supportedProtocols: (JVM default)
      excludedProtocols: [SSL, SSLv2, SSLv2Hello, SSLv3] # (Jetty's default)
      supportedCipherSuites: (JVM default)
      excludedCipherSuites: [.*_(MD5|SHA|SHA1)$] # (Jetty's default)
      allowRenegotiation: true
      endpointIdentificationAlgorithm: (none)
      disableSniHostCheck: false
```

Name	Default	Description
key-StorePath	REQUIRED	The path to the Java key store which contains the host certificate and private key.
key-StorePassword	REQUIRED	The password used to access the key store.
key-Store-Type	JKS	The type of key store (usually JKS, PKCS12, JCEKS, Windows-MY, or Windows-ROOT).
key-Store-Provider	(none)	The JCE provider to use to access the key store.
trust-StorePath	(none)	The path to the Java key store which contains the CA certificates used to establish trust.
trust-StorePassword	(none)	The password used to access the trust store.
trust-Store-Type	JKS	The type of trust store (usually JKS, PKCS12, JCEKS, Windows-MY, or Windows-ROOT).
trust-Store-Provider	(none)	The JCE provider to use to access the trust store.
key-Manager-Password	(none)	The password, if any, for the key manager.
need-ClientAuth	(none)	Whether or not client authentication is required.
want-ClientAuth	(none)	Whether or not client authentication is requested.
certAlias	(none)	The alias of the certificate to use.
crlPath	(none)	The path to the file which contains the Certificate Revocation List.
enable-CRLDP	false	Whether or not CRL Distribution Points (CRLDP) support is enabled.
enableOCSP	false	Whether or not On-Line Certificate Status Protocol (OCSP) support is enabled.
max-Cert-Path-Length	(unlimited)	The maximum certification path length.
ocspResponderUrl	(none)	The location of the OCSP responder.
jce-Provider	(none)	The name of the JCE provider to use for cryptographic support. See Oracle documentation for more information.
validate-Certs	false	Whether or not to validate TLS certificates before starting. If enabled, Dropwizard will refuse to start with expired or otherwise invalid certificates. This option will cause unconditional failure in Dropwizard 1.x until a new validation mechanism can be implemented.
validate-Peers	false	Whether or not to validate TLS peer certificates. This option will cause unconditional failure in Dropwizard 1.x until a new validation mechanism can be implemented.

18.2.3 HTTP/2 over TLS

HTTP/2 is a new protocol, intended as a successor of HTTP/1.1. It adds several important features like binary structure, stream multiplexing over a single connection, header compression, and server push. At the same time it remains semantically compatible with HTTP/1.1, which should make the upgrade process more seamless. Checkout [HTTP/2 FAQ](#) for the further information.

For an encrypted connection HTTP/2 uses ALPN protocol. It's a TLS extension, that allows a client to negotiate a protocol to use after the handshake is complete. If either side does not support ALPN, then the protocol will be ignored, and an HTTP/1.1 connection over TLS will be used instead.

For this connector to work with ALPN protocol you need to either:

- Enable native SSL support via Google's Conscrypt as described in the [SSL section](#) of the Core manual; or
- Provide alpn-boot library to JVM's bootpath. The correct library version depends on the JVM version. Consult [Jetty ALPN guide](#) for the reference.

Note that your JVM also must provide `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256` cipher. The specification [states](#) that HTTP/2 deployments must support it to avoid handshake failures. It's the single supported cipher in HTTP/2 connector by default. In case you want to support more strong ciphers, you should specify them in the `supportedCipherSuites` parameter along with `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`.

This connector extends the attributes that are available to the [HTTPS connector](#)

```
server:
  applicationConnectors:
    - type: h2
      port: 8445
      maxConcurrentStreams: 1024
      initialStreamRecvWindow: 65535
      keyStorePath: /path/to/file # required
      keyStorePassword: changeit
      trustStorePath: /path/to/file # required
      trustStorePassword: changeit
      supportedCipherSuites: # optional
        - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
        - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
```

Name	De- fault	Description
maxConcurrentStreams	1024	The maximum number of concurrently open streams allowed on a single HTTP/2 connection. Larger values increase parallelism, but cost a memory commitment.
initialStreamRecvWindow	65535	The initial flow control window size for a new stream. Larger values may allow greater throughput, but also risk head of line blocking if TCP/IP flow control is triggered.

18.2.4 HTTP/2 Plain Text

HTTP/2 promotes using encryption, but doesn't require it. However, most browsers stated that they will not support HTTP/2 without encryption. Currently no browser supports HTTP/2 unencrypted.

The connector should only be used in closed secured networks or during development. It expects from clients an HTTP/1.1 OPTIONS request with Upgrade : h2c header to indicate a wish to upgrade to HTTP/2, or a request with the HTTP/2 connection preface. If the client doesn't support HTTP/2, a plain HTTP/1.1 connections will be used instead.

This connector extends the attributes that are available to the *HTTP connector*

```
server:
  applicationConnectors:
    - type: h2c
      port: 8446
      maxConcurrentStreams: 1024
      initialStreamRecvWindow: 65535
```

Name	De- fault	Description
maxConcurrentStreams	1024	The maximum number of concurrently open streams allowed on a single HTTP/2 connection. Larger values increase parallelism, but cost a memory commitment.
initialStreamRecvWindow	65535	The initial flow control window size for a new stream. Larger values may allow greater throughput, but also risk head of line blocking if TCP/IP flow control is triggered.

18.3 Tasks

```
admin:
  tasks:
    printStackTraceOnError: true
```

Name	Default	Description
printStackTraceOnError	false	Print the full stack trace when the execution of a task failed.

18.4 Health checks

```
admin:
  healthChecks:
    servletEnabled: true
    minThreads: 1
    maxThreads: 4
    workQueueSize: 1
```

Name	Default	Default Description
servletEnabled	true	Whether to enable or disable the health check servlet.
minThreads	1	The minimum number of threads for executing health checks.
maxThreads	4	The maximum number of threads for executing health checks.
workQueueSize	1	The length of the work queue for health check executions.

18.5 Logging

```
logging:
  level: INFO
  loggers:
    "io.dropwizard": INFO
    "org.hibernate.SQL":
      level: DEBUG
      additive: false
      appenders:
        - type: file
          currentLogFilename: /var/log/myapplication-sql.log
          archivedLogFilenamePattern: /var/log/myapplication-sql-%d.log.gz
          archivedFileCount: 5
  appenders:
    - type: console
```

Name	Default	Description
level	Level.INFO	Logback logging level.
additive	true	Logback additive setting.
loggers	(none)	Individual logger configuration (both forms are acceptable).
appenders	(none)	One of console, file or syslog.

18.5.1 Console

```
logging:
  level: INFO
  appenders:
    - type: console
      threshold: ALL
      queueSize: 512
      discardingThreshold: 0
      timeZone: UTC
      target: stdout
      logFormat: "%-5p [%d{ISO8601,UTC}] %c: %m%n%dwREx"
      filterFactories:
        - type: URI
```

Name	Default	Description
type	REQUIRED	The appender type. Must be <code>console</code> .
threshold	ALL	The lowest level of events to print to the console.
queueSize	256	The maximum capacity of the blocking queue.
discardingThreshold	-1	When the blocking queue has only the capacity mentioned in <code>discardingThreshold</code> remaining, it will drop events of level <code>TRACE</code> , <code>DEBUG</code> and <code>INFO</code> , keeping only events of level <code>WARN</code> and <code>ERROR</code> . If no discarding threshold is specified (-1), then a default of <code>queueSize / 5</code> (logback's default ratio) is used. To keep all events, set <code>discardingThreshold</code> to 0.
timeZone	UTC	The time zone to which event timestamps will be converted. To use the system/default time zone, set it to <code>system</code> .
target	stdout	The name of the standard stream to which events will be written. Can be <code>stdout</code> or <code>stderr</code> .
logFormat	<code>%-5p [%d{ISO8601,UTC}] %c: %m%n%dwREx</code>	The Logback pattern with which events will be formatted. See the Logback documentation for details. The default log pattern is <code>%h %l %u [%t{dd/MMM/yyyy:HH:mm:ss Z,UTC}] "%r" %s %b "%i{Referer}" "%i{User-Agent}" %D`</code> . Use the placeholder <code>%m%n%dwREx</code> to include the value of <code>timeZone</code> in the pattern.
filterFactories	(none)	The list of filters to apply to the appender, in order, after the threshold.
neverBlock	false	Prevent the wrapping asynchronous appender from blocking when its underlying queue is full. Set to true to disable blocking.

18.5.2 File

```

logging:
  level: INFO
  appenders:
    - type: file
      currentLogFilename: /var/log/myapplication.log
      threshold: ALL
      queueSize: 512
      discardingThreshold: 0
      archive: true
      archivedLogFilenamePattern: /var/log/myapplication-%d.log
      archivedFileCount: 5
      timeZone: UTC
      logFormat: "%-5p [%d{ISO8601,UTC}] %c: %m%n%dwREx"
      bufferSize: 8KiB
      immediateFlush: true
      filterFactories:
        - type: URI

```

Name	Default	Description
type	REQUIRED	The appender type. Must be <code>file</code> .
current-Log-File-name	REQUIRED	The filename where current events are logged.
threshold	ALL	The lowest level of events to write to the file.
queue-Size	256	The maximum capacity of the blocking queue.
discardingThreshold	-1	When the blocking queue has only the capacity mentioned in <code>discardingThreshold</code> remaining, it will drop events of level TRACE, DEBUG and INFO, keeping only events of level WARN and ERROR. If no discarding threshold is specified (-1), then a default of <code>queueSize / 5</code> (logback's default ratio) is used. To keep all events, set <code>discardingThreshold</code> to 0.
archive	true	Whether or not to archive old events in separate files.
archived-Log-File-namePattern	(none)	Required if <code>archive</code> is true. The filename pattern for archived files. If <code>maxFileSize</code> is specified, rollover is size-based, and the pattern must contain <code>%i</code> for an integer index of the archived file. Otherwise rollover is date-based, and the pattern must contain <code>%d</code> , which is replaced with the date in <code>yyyy-MM-dd</code> form. If the pattern ends with <code>.gz</code> or <code>.zip</code> , files will be compressed as they are archived.
archived-File-Count	5	The number of archived files to keep. Must be greater than or equal to 0. Zero is a special value signifying to keep infinite logs (use with caution)
max-File-Size	(unlimited)	The maximum size of the currently active file before a rollover is triggered. The value can be expressed in bytes, kibibytes, kilobytes, mebibytes, megabytes, gibibytes, gigabytes, tebibytes, terabytes, pebibytes, and petabytes by appending B, KiB, KB, MiB, MB, GiB, GB, TiB, TB, PiB, or PB to the numeric value. Examples include 5KiB, 100MiB, 1GiB, 1TB. Sizes can also be spelled out, such as 5 kibibytes, 100 mebibytes, 1 gibibyte, 1 terabyte.
total-Size-Cap	(unlimited)	Controls the total size of all files. Oldest archives are deleted asynchronously when the total size cap is exceeded.
time-Zone	UTC	The time zone to which event timestamps will be converted.
log-Format	<code>%-5p [%d{ISO8601}UTC] [%c: %m%n%dw%n]</code>	The Logback pattern with which events will be formatted. See the Logback documentation for details. The default log pattern is <code>`%h %l %u [%t{dd/MMM/yyyy:HH:mm:ss Z,UTC}] "%r" %s %b "%i{Referer}" "%i{User-Agent}" %D`</code> . Use the placeholder <code>%m%n%dw%n]</code> to include the value of <code>timeZone</code> in the pattern.
filter-Factories	(none)	The list of filters to apply to the appender, in order, after the threshold.
neverBlock	false	Prevent the wrapping asynchronous appender from blocking when its underlying queue is full. Set to true to disable blocking.
buffer-Size	8KiB	The buffer size of the underlying FileAppender (setting added in logback 1.1.10). Increasing this from the default of 8KiB to 256KiB is reported to significantly reduce thread contention.
immediate-Flush	true	If set to true, log events will be immediately flushed to disk. Immediate flushing is safer, but it degrades logging throughput.

18.5.3 Syslog

```

logging:
  level: INFO
  appenders:
    - type: syslog
      host: localhost
      port: 514
      facility: local0
      threshold: ALL
      stackTracePrefix: \t
      logFormat: "%-5p [%d{ISO8601,UTC}] %c: %m%n%dwREx"
      filterFactories:
        - type: URI

```

Name	Default	Description
host	localhost	The hostname of the syslog server.
port	514	The port on which the syslog server is listening.
facility	local0	The syslog facility to use. Can be either auth, authpriv, daemon, cron, ftp, lpr, kern, mail, news, syslog, user, uucp, local0, local1, local2, local3, local4, local5, local6, or local7.
threshold	ALL	The lowest level of events to write to the file.
logFormat	%-5p [%d{ISO8601,UTC}] %c: %m%n%dwREx	The Logback pattern with which events will be formatted. See the Logback documentation for details. The default log pattern is <code>%h %l %u [%t{dd/MMM/yyyy:HH:mm:ss Z,UTC}] "%r" %s %b "%i{Referer}" "%i{User-Agent}" %D`</code> .
stackTracePrefix	t	The prefix to use when writing stack trace lines (these are sent to the syslog server separately from the main message)
filterFactories	(none)	The list of filters to apply to the appender, in order, after the threshold.
neverBlock	false	Prevent the wrapping asynchronous appender from blocking when its underlying queue is full. Set to <code>true</code> to disable blocking.
includeStackTrace	true	Set to <code>false</code> to disable sending stack traces to the syslog service.

18.5.4 TCP

```

logging:
  level: INFO
  appenders:
    - type: tcp
      host: localhost
      port: 4560
      connectionTimeout: 500ms

```

(continues on next page)

(continued from previous page)

```

immediateFlush: true
sendBufferSize: 8KiB

```

Name	De- fault	Description
host	local- host	The hostname of the TCP server.
port	4560	The port on which the TCP server is listening.
connection- Timeout	500ms	The timeout to connect to the TCP server.
immediate- Flush	true	If set to true, log events will be immediately send to the server Immediate flushing is safer, but it degrades logging throughput.
sendBuffer- Size	8KiB	The buffer size of the underlying SocketAppender. Takes into effect if immediateFlush is disabled.

18.5.5 UDP

```

logging:
  level: INFO
  appenders:
    - type: udp
      host: localhost
      port: 514

```

Name	Default	Description
host	localhost	The hostname of the UDP server.
port	514	The port on which the UDP server is listening.

18.5.6 FilterFactories

A factory used for request logging appenders should implement `io.dropwizard.logging.common.filter.FilterFactory<IAccessEvent>` while one used for regular logging should implement `io.dropwizard.logging.common.filter.FilterFactory<ILoggingEvent>`. To register a factory, its fully qualified classname must be listed in `META-INF/services/io.dropwizard.logging.common.filter.FilterFactory`. The factory then can be referenced in the configuration either via its simple classname or via type name, if factory class annotated with `@JsonTypeName`.

```

logging:
  level: INFO
  appenders:
    - type: console
      filterFactories:
        - type: URI

```

Name	Default	Description
type	REQUIRED	The filter type name.

18.5.7 JSON layout

```
layout:
  type: json
  timestampFormat: "yyyy-MM-dd'T'HH:mm:ss.SSSZ"
  prettyPrint: false
  appendLineSeparator: true
  includes: [timestamp, threadName, level, loggerName, message, mdc, exception]
  customFieldNames:
    timestamp: "@timestamp"
  additionalFields:
    service-name: "user-service"
  includesMdcKeys: [userId]
  flattenMdc: true
  exception:
    rootFirst: true
    depth: full
    evaluators: [org.apache]
```

Name	Default	Description
timestampFormat	(none)	By default, the timestamp is not formatted. To customize how timestamps are formatted, set the property to the corresponding <code>DateTimeFormatter</code> string or one of the predefined formats (e.g. <code>ISO_LOCAL_TIME</code> , <code>ISO_ZONED_DATE_TIME</code> , <code>RFC_1123_DATE_TIME</code>).
prettyPrint	false	Whether the JSON output should be formatted for human readability.
appendLineSeparator	true	Whether to append a line separator at the end of the message formatted as JSON.
includes	(timestamp, level, threadName, mdc, loggerName, message, exception)	Set of logging event attributes to include in the JSON map: <ul style="list-style-type: none"> • <code>timestamp true</code> Whether to include the timestamp as the <code>timestamp</code> field. • <code>level true</code> Whether to include the logging level as the <code>level</code> field. • <code>threadName true</code> Whether to include the thread name as the <code>thread</code> field. • <code>mdc true</code> Whether to include the MDC properties as the <code>mdc</code> field. • <code>loggerName true</code> Whether to include the logger name as the <code>logger</code> field. • <code>message true</code> Whether to include the formatted message as the <code>message</code> field. • <code>exception true</code> Whether to log exceptions. If the property enabled and there is an exception, it will be formatted to a string as the <code>exception</code> field. • <code>contextName false</code> Whether to include the logging context name as the <code>context</code> field.
customFieldNames	(empty)	Map of field name replacements. For example (<code>requestTime:request_time</code> , <code>userAgent:user_agent</code>).
additionalFields	(empty)	Map of fields to add in the JSON map.
includesMdcKeys	(empty)	Set of MDC keys which should be included in the JSON map. By default includes everything.
flattenMdc	false	Flatten the MDC to the root of the JSON object instead of nested in the
164	Chapter 18. Dropwizard Configuration Reference	
exception	(empty)	The <code>exception</code> configuration for the <code>exception</code> field.

Exception

```

layout:
  type: json
  exception:
    rootFirst: false
    depth: 25
    evaluators: [org.apache]

```

Name	Default	Description
rootFirst	true	Whether the root cause should be displayed first.
depth	full	The stack trace <code>depth</code> .
evaluators	(empty)	The packages to <code>filter</code> from the stacktrace.

18.5.8 JSON access log layout

```

layout:
  type: access-json
  timestampFormat: "yyyy-MM-dd'T'HH:mm:ss.SSSZ"
  prettyPrint: false
  appendLineSeparator: true
  includes: [timestamp, remoteAddress, remoteUser, protocol, method, requestUri,
↳statusCode, requestTime, contentLength, userAgent]
  requestHeaders:
    - X-Request-Id
  responseHeaders:
    - X-Request-Id
  requestAttributes:
    - SomeAttributeName
  customFieldNames:
    timestamp: "@timestamp"
  additionalFields:
    service-name: "user-service"

```

Name	Default	Description
timestampFormat	(none)	By default, the timestamp is not formatted. To customize how timestamps are formatted, set the property to the corresponding <code>DateTimeFormatter</code> string or one of the predefined formats (e.g. <code>ISO_LOCAL_TIME</code> , <code>ISO_ZONED_DATE_TIME</code> , <code>``RFC_1123_DATE_TIME``</code>)
prettyPrint	false	Whether the JSON output should be formatted for human readability.
appendLineSeparator	true	Whether to append a line separator at the end of the message formatted as JSON.
includes	(timestamp, remoteAddress, protocol, method, requestUri, statusCode, requestTime, contentLength, userAgent)	<p>Set of logging event attributes to include in the JSON map:</p> <ul style="list-style-type: none"> • <code>contentLength</code> <i>true</i> Whether to include the response content length, if it's known as the <code>contentLength</code> field. • <code>method</code> <i>true</i> Whether to include the request HTTP method as the <code>method</code> field. • <code>remoteAddress</code> <i>true</i> Whether to include the IP address of the client or last proxy that sent the request as the <code>remoteAddress</code> field. • <code>remoteUser</code> <i>true</i> Whether to include information about the remote user as the <code>remoteUser</code> field. • <code>requestTime</code> <i>true</i> Whether to include the time elapsed between receiving the request and logging it as the <code>requestTime</code> field. Time is in <i>ms</i>. • <code>requestUri</code> <i>true</i> Whether to include the URI of the request as the <code>uri</code> field. • <code>statusCode</code> <i>true</i> Whether to include the status code of the response as the <code>status</code> field. • <code>protocol</code> <i>true</i> Whether to include the request HTTP protocol as the <code>protocol</code> field. • <code>timestamp</code> <i>true</i> Whether to include the timestamp of the event the <code>timestamp</code> field. • <code>userAgent</code> <i>true</i> Whether to include the user agent of the request as the <code>userAgent</code> field. • <code>requestParameters</code> <i>false</i>

18.6 Metrics

The metrics configuration has three fields; frequency, reporters and reportOnStop.

```
metrics:
  frequency: 1 minute
  reporters:
    - type: <type>
  reportOnStop: false
```

Name	Default	Description
frequency	1 minute	The frequency to report metrics. Overridable per-reporter.
reporters	(none)	A list of reporters to report metrics.
reportOnStop	false	To report metrics one last time when stopping Dropwizard.

18.6.1 All Reporters

The following options are available for all metrics reporters.

```
metrics:
  reporters:
    - type: <type>
      durationUnit: milliseconds
      rateUnit: seconds
      excludes: (none)
      includes: (all)
      excludesAttributes: (none)
      includesAttributes: (all)
      useRegexFilters: false
      frequency: 1 minute
```

Name	Default	Description
durationUnit	milliseconds	The unit to report durations as. Overrides per-metric duration units.
rateUnit	seconds	The unit to report rates as. Overrides per-metric rate units.
excludes	(none)	Metrics to exclude from reports, by name. When defined, matching metrics will not be reported.
includes	(all)	Metrics to include in reports, by name. When defined, only these metrics will be reported.
excludesAttributes	(none)	Metric attributes to exclude from reports, by name (e.g. p98, m15_rate, stddev). When defined, matching metrics attributes will not be reported.
includesAttributes	(all)	Metrics attributes to include in reports, by name (e.g. p98, m15_rate, stddev). When defined, only these attributes will be reported.
useRegexFilters	false	Indicates whether the values of the 'includes' and 'excludes' fields should be treated as regular expressions or not.
useSubstringMatching	false	Uses a substring matching strategy to determine whether a metric should be processed.
frequency	(none)	The frequency to report metrics. Overrides the default.

The inclusion and exclusion rules are defined as:

- If **includes** is empty, then all metrics are included;
- If **includes** is not empty, only metrics from this list are included;
- If **excludes** is empty, no metrics are excluded;
- If **excludes** is not empty, then exclusion rules take precedence over inclusion rules. Thus if a name matches the exclusion rules it will not be included in reports even if it also matches the inclusion rules.

When neither **useRegexFilters** nor **useSubstringMatching** are enabled, a default exact matching strategy will be used to determine whether a metric should be processed. In case both **useRegexFilters** and **useSubstringMatching** are set, **useRegexFilters** takes precedence over **useSubstringMatching**.

Formatted Reporters

These options are available only to “formatted” reporters and extend the options available to *all reporters*

```
metrics:
  reporters:
    - type: <type>
      locale: <system default>
```

Name	Default	Description
locale	System default	The Locale for formatting numbers, dates and times.

18.6.2 Console Reporter

Reports metrics periodically to the console.

Extends the attributes that are available to *formatted reporters*

```
metrics:
  reporters:
    - type: console
      timeZone: UTC
      output: stdout
```

Name	Default	Description
timeZone	UTC	The timezone to display dates/times for.
output	stdout	The stream to write to. One of stdout or stderr.

18.6.3 CSV Reporter

Reports metrics periodically to a CSV file.

Extends the attributes that are available to *formatted reporters*

```
metrics:
  reporters:
    - type: csv
      file: /path/to/file
```

Name	Default	Description
file	No default	The CSV file to write metrics to.

18.6.4 Graphite Reporter

Reports metrics periodically to Graphite.

Extends the attributes that are available to *all reporters*

Note: You will need to add `dropwizard-metrics-graphite` to your POM.

```
metrics:
  reporters:
    - type: graphite
      host: localhost
      port: 2003
      prefix: <prefix>
      transport: tcp
```

Name	Default	Description
host	localhost	The hostname of the Graphite server to report to.
port	2003	The port of the Graphite server to report to.
prefix	(none)	The prefix for Metric key names to report to Graphite.
transport	tcp	The type of transport to report to Graphite with ("tcp" or "udp").

18.6.5 SLF4J

Reports metrics periodically by logging via SLF4J.

Extends the attributes that are available to *all reporters*

See [BaseReporterFactory](#) and [BaseFormattedReporterFactory](#) for more options.

```
metrics:  
  reporters:  
    - type: log  
      logger: metrics  
      markerName: <marker name>
```

Name	Default	Description
logger	metrics	The name of the logger to write metrics to.
markerName	(none)	The name of the marker to mark logged metrics with.

18.7 Health

```
health:  
  enabled: true  
  delayedShutdownHandlerEnabled: true  
  shutdownWaitPeriod: 5s  
  healthCheckUrlPaths: ["/health-check"]  
  healthChecks:  
    - <some health check config>  
    - <some other health check config>  
  initialOverallState: false  
  responseProvider:  
    type: json  
  responder:  
    type: servlet
```

Name	Default	Description
enabled	true	Flag indicating whether to enable health functionality or not.
delayed- Shutdown- HandlerEn- abled	false	Flag indicating whether to delay shutdown to allow already processing requests to complete.
shutdown- WaitPeriod	15 sec- onds	Amount of time to delay shutdown by to allow already processing requests to complete. Only applicable if <code>delayedShutdownHandlerEnabled</code> is true.
healthCheck- UrlPaths	["/health- check"]	URLs to expose the app's health check on.
healthChecks	[]	A list of configured health checks. See the [Health Check Configuration section](#health-check-configuration) for more details.
initialOver- allState	true	Flag indicating whether the overall health state of the application should start as healthy or unhealthy. A value of <code>true</code> indicates an initial state of healthy while a value of <code>false</code> indicates an initial state of unhealthy.
response- Provider	json	The health response provider that is used to respond to generate responses to return to health check requests. This can be implemented using Jersey, Jetty, or other technologies if desired. See the detailed JSON health response provider section for more details.
responder	servlet	The health responder that is used to respond to health check requests. This can be implemented using Jersey, Jetty, or other technologies if desired. See the servlet health responder section for more details.

18.7.1 Health Checks

Options around a particular health check which is registered in an Application

```
health:
  healthChecks:
    - name: file-system
      type: alive
      critical: true
      initialState: true
    - name: database
      type: ready
      critical: false
      initialState: false
```

Name	Default	Description
name	(none)	The name of this health check. This must be unique, and match the name of the check registered in code. (On the application's <code>HealthCheckRegistry</code>)
type	ready	The type of this health check. This is either <code>alive</code> or <code>ready</code> . See the application status section for more details.
critical	false	Flag indicating whether this dependency is critical to determine the health of the application. If <code>true</code> and this dependency is unhealthy, the application will also be marked as unhealthy.
initial-State	true	Flag indicating the initial state to use for this health check. A value of <code>true</code> indicates an initial state of healthy while a value of <code>false</code> indicates an initial state of unhealthy.
schedule	default schedule	The schedule that this health check will be run on. See the schedule section for more details.

18.7.2 Schedule

The schedule on which to execute a particular *health checks*

```

health:
  healthChecks:
    - name: file-system
      schedule:
        checkInterval: 10s
        downtimeInterval: 2s
        initialDelay: 5s
        failureAttempts: 1
        successAttempts: 2

```

Name	Default	Description
checkInterval	5 seconds	The interval on which to perform a health check for this dependency while the dependency is in a healthy state.
downtimeInterval	30 seconds	The interval on which to perform a health check for this dependency while the dependency is in an unhealthy state.
initialDelay	the value of <i>checkInterval</i>	The initial delay to use when first scheduling the health check.
failureAttempts	3	The threshold of consecutive failed attempts needed to mark a dependency as unhealthy (from a healthy state).
successAttempts	2	The threshold of consecutive successful attempts needed to mark a dependency as healthy (from an unhealthy state).

18.7.3 Detailed JSON Health Response Provider

A detailed servlet used to handle health check requests, which returns a JSON response explaining the various registered health checks, their current status, and other metadata.

```
health:
  responseProvider:
    type: json
```

Name	Default	Description
type	json - json	

18.7.4 Servlet Health Responder

A servlet responder used to handle health check requests.

```
health:
  responder:
    type: servlet
    cacheControlEnabled: true
    cacheControlValue: "no-store"
```

Name	Default	Description
type	servlet	<ul style="list-style-type: none"> servlet
cacheControlEnabled	true	Flag controlling whether a Cache-Control header will be included in the health check response or not. Set header value using cacheControlValue.
cacheControlValue	"no-store"	The value to be set in the Cache-Control header in the health check response. Only used if cacheControlEnabled is set to true.

18.8 Clients

18.8.1 HttpClient

See [HttpClientConfiguration](#) for more options.

```
httpClient:
  timeout: 500ms
  connectionTimeout: 500ms
  timeToLive: 1h
  cookiesEnabled: false
  maxConnections: 1024
```

(continues on next page)

(continued from previous page)

```

maxConnectionsPerRoute: 1024
keepAlive: 0ms
retries: 0
userAgent: <application name> (<client name>)

```

Name	Default	Description
timeout	500 milliseconds	The maximum idle time for a connection, once established.
connectionTimeout	500 milliseconds	The maximum time to wait for a connection to open.
connectionRequestTimeout	500 milliseconds	The maximum time to wait for a connection to be returned from the connection pool.
timeToLive	1 hour	The maximum time a pooled connection can stay idle (not leased to any thread) before it is shut down.
cookiesEnabled	false	Whether or not to enable cookies.
maxConnections	1024	The maximum number of concurrent open connections.
maxConnectionsPerRoute	1024	The maximum number of concurrent open connections per route.
keepAlive	0 milliseconds	The maximum time a connection will be kept alive before it is reconnected. If set to 0, connections will be immediately closed after every request/response.
retries	0	The number of times to retry failed requests. Requests are only retried if they throw an exception other than <code>InterruptedException</code> , <code>UnknownHostException</code> , <code>ConnectException</code> , or <code>SSLException</code> .
userAgent	applicationName (clientName)	The User-Agent to send with requests.
validateAfterInactivityPeriod	0 milliseconds	The maximum time before a persistent connection is checked to remain active. If set to 0, no inactivity check will be performed.

Proxy

```

httpClient:
  proxy:
    host: 192.168.52.11
    port: 8080
    scheme : http
    auth:
      username: secret
      password: stuff
      authScheme: NTLM
      realm: realm
      hostname: host
      domain: WINDOWSDOMAIN
      credentialType: NT

```

(continues on next page)

(continued from previous page)

nonProxyHosts:

- localhost
- '192.168.52.*'
- '*.example.com'

Name	Default	Description
host	REQUIRED	The proxy server host name or ip address.
port	(scheme default)	The proxy server port. If the port is not set then the scheme default port is used.
scheme	http	The proxy server URI scheme. HTTP and HTTPS schemas are permitted. By default HTTP scheme is used.
auth	(none)	The proxy server Basic or NTLM authentication schemes. If they are not set then no credentials will be passed to the server.
username	REQUIRED	The username used to connect to the server.
password	REQUIRED	The password used to connect to the server.
authScheme	Basic	The authentication scheme used by the. Allowed options are: Basic, NTLM
realm	(none)	The realm, used for NTLM authentication.
hostname	(none)	The hostname of the windows workstation, used for NTLM authentication.
domain	(none)	The Windows Domain, used for NTLM authentication.
credentialType	(none)	The Apache HTTP Client Credentials implementation used for proxy authentication. Allowed options are: UsernamePassword or NT
nonProxyHosts	(none)	List of patterns of hosts that should be reached without proxy. The patterns may contain symbol '*' as a wildcard. If a host matches one of the patterns it will be reached through a direct connection.

TLS**httpClient:****tls:**

```

protocol: TLSv1.2
provider: SunJSSE
verifyHostname: true
keyStorePath: /path/to/file
keyStorePassword: changeit
keyStoreType: JKS
trustStorePath: /path/to/file
trustStorePassword: changeit
trustStoreType: JKS
trustSelfSignedCertificates: false
supportedProtocols: TLSv1.1,TLSv1.2

```

(continues on next page)

(continued from previous page)

```
supportedCipherSuites: TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
certAlias: alias-of-specific-cert
```

Name	Default	Description
protocol	TLSv1	The default protocol the client will attempt to use during the SSL Handshake. See here for more information.
provider	(none)	The name of the JCE provider to use on client side for cryptographic support (for example, SunJCE, Conscrypt, BC, etc). See Oracle documentation for more information.
verifyHostName	true	Whether to verify the hostname of the server against the hostname presented in the server certificate.
keyStorePath	(none)	The path to the Java key store which contains the client certificate and private key.
keyStorePassword	(none)	The password used to access the key store.
keyStoreType	JKS	The type of key store (usually JKS, PKCS12, JCEKS, Windows-MY, or Windows-ROOT).
trustStorePath	(none)	The path to the Java key store which contains the CA certificates used to establish trust.
trustStorePassword	(none)	The password used to access the trust store.
trustStoreType	JKS	The type of trust store (usually JKS, PKCS12, JCEKS, Windows-MY, or Windows-ROOT).
trustSelfSignedCertificates	false	If true, will trust all self-signed certificates regardless of trustStore settings. If false, trust decisions will be handled by the supplied trustStore.
supportedProtocols	(none)	A list of protocols (e.g., SSLv3, TLSv1) which are supported. All other protocols will be refused.
supportedCipherSuites	(none)	A list of cipher suites (e.g., TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256) which are supported. All other cipher suites will be refused.
certAlias	(none)	The alias of a specific client certificate to present when authenticating. Use this when the specified keystore has multiple certificates to force use of a non-default certificate.

18.8.2 JerseyClient

Extends the attributes that are available to *http clients*

See [JerseyClientConfiguration](#) and [HttpClientConfiguration](#) for more options.

```
jerseyClient:
  minThreads: 1
  maxThreads: 128
  workQueueSize: 8
  gzipEnabled: true
  gzipEnabledForRequests: true
  chunkedEncodingEnabled: true
```

Name	Default	Description
minThreads	5	The minimum number of threads in the pool used for asynchronous requests.
max-Threads	128	The maximum number of threads in the pool used for asynchronous requests. If asynchronous requests made by jersey client while serving requests, the number must be set according to the <i>maxThread</i> setting of the <i>server</i> . Otherwise some requests made to dropwizard on heavy load may fail due to congestion on the jersey client's thread pool.
workQueueSize	5	The size of the work queue of the pool used for asynchronous requests. Additional threads will be spawned only if the queue is reached its maximum size.
gzipEnabled	true	Adds an Accept-Encoding: gzip header to all requests, and enables automatic gzip decoding of responses.
gzipEnabled-ForRequests	true	Adds a Content-Encoding: gzip header to all requests, and enables automatic gzip encoding of requests.
chunkedEncodingEnabled	true	Enables the use of chunked encoding for requests.

18.9 Database

```

database:
  driverClass : org.postgresql.Driver
  url: 'jdbc:postgresql://db.example.com/db-prod'
  user: pg-user
  password: iAMs00perSecrEET

```

Name	Default	Description
url	REQUIRED	The URL of the server.
driverClass	none	The fully qualified class name of the JDBC driver class. Only required if the driver is not in the classpath.
user	none	The username used to connect to the server.
password	none	The password used to connect to the server.
removeAbandoned	false	Remove abandoned connections if they exceed removeAbandonedTimeout. If true, the connections that have been abandoned (timed out) won't get closed and reported as abandoned.
removeAbandonedTimeout	60 seconds	The time before a database connection can be considered abandoned.
abandonWhenPercentageFull	0	Connections that have been abandoned (timed out) won't get closed and reported as abandoned.
alternateUsernamesAllowed	false	Set to true if the call getConnection(username,password) is allowed. This is useful for applications that use a single connection pool to connect to multiple servers.
commitOnReturn	false	Set to true if you want the connection pool to commit any pending transactions when a connection is returned to the pool.
rollbackOnReturn	false	Set to true if you want the connection pool to rollback any pending transactions when a connection is returned to the pool.
autoCommitByDefault	JDBC driver's default	The default auto-commit state of the connections.
readOnlyByDefault	JDBC driver's default	The default read-only state of the connections.
properties	none	Any additional JDBC driver parameters.
defaultCatalog	none	The default catalog to use for the connections.
defaultTransactionIsolation	JDBC driver's default	The default transaction isolation to use for the connections. Can be one of none, read committed, repeatable read, or serializable.
useFairQueue	true	If true, calls to getConnection are handled in a FIFO manner.
initialSize	10	The initial size of the connection pool.
minSize	10	The minimum size of the connection pool.

Name	Default	Description
maxSize	100	The maximum size of the connection pool.
initializationQuery	none	A custom query to be run when a connection is first created.
logAbandonedConnections	false	If true, logs stack traces of abandoned connections.
logValidationErrors	false	If true, logs errors when connections fail validation.
maxConnectionAge	none	If set, connections which have been open for longer than maxConnectionAge
maxWaitForConnection	30 seconds	If a request for a connection is blocked for longer than this period, an excepti
minIdleTime	1 minute	The minimum amount of time an connection must sit idle in the pool before
validationQuery	SELECT 1	The SQL query that will be used to validate connections from this pool before
validationQueryTimeout	none	The timeout before a connection validation queries fail.
checkConnectionWhileIdle	true	Set to true if query validation should take place while the connection is idle.
checkConnectionOnBorrow	false	Whether or not connections will be validated before being borrowed from the
checkConnectionOnConnect	false	Whether or not connections will be validated before being added to the pool.
checkConnectionOnReturn	false	Whether or not connections will be validated after being returned to the pool.
autoCommentsEnabled	true	Whether or not ORMs should automatically add comments.
evictionInterval	5 seconds	The amount of time to sleep between runs of the idle connection validation, a
validationInterval	30 seconds	To avoid excess validation, only run validation once every interval.
validatorClassName	none	Name of a class of a custom validator implementation, which will be used fo
jdbcInterceptors	none	A semicolon separated list of JDBC interceptor classnames.
ignoreExceptionOnPreLoad	false	Flag whether ignore error of connection creation while initializing the pool.

18.10 Polymorphic configuration

The `dropwizard-configuration` module provides you with a polymorphic configuration mechanism, meaning that a particular section of your configuration file can be implemented using one or more configuration classes.

To use this capability for your own configuration classes, create a top-level configuration interface or class that implements `Discoverable` and add the name of that class to `META-INF/services/io.dropwizard.jackson.Discoverable`. Make sure to use [Jackson polymorphic deserialization](#) annotations appropriately.

```
@JsonTypeInfo(use = Id.NAME, include = As.PROPERTY, property = "type")
interface WidgetFactory extends Discoverable {
    Widget createWidget();
}
```

Then create subtypes of the top-level type corresponding to each alternative, and add their names to `META-INF/services/WidgetFactory`.

```
@JsonTypeName("hammer")
public class HammerFactory implements WidgetFactory {
    @JsonProperty
    private int weight = 10;

    @Override
    public Hammer createWidget() {
        return new Hammer(weight);
    }
}
```

(continues on next page)

(continued from previous page)

```
@JsonTypeName("chisel")
public class ChiselFactory implements WidgetFactory {
    @JsonProperty
    private float radius = 1;

    @Override
    public Chisel createWidget() {
        return new Chisel(radius);
    }
}
```

Now you can use `WidgetFactory` objects in your application's configuration.

```
public class MyConfiguration extends Configuration {
    @JsonProperty
    @NotNull
    @Valid
    private List<WidgetFactory> widgets;
}
```

```
widgets:
- type: hammer
  weight: 20
- type: chisel
  radius: 0.4
```

See *testing configurations* for details on ensuring the configuration will be deserialized correctly.

DROPWIZARD INTERNALS

You already read through the whole Dropwizard documentation? Congrats! Then you are ready to have a look into some nitty-gritty details of Dropwizard.

19.1 Startup Sequence

`Application<T>` extends `Configuration` is the “Main” class of a Dropwizard Application.

`application.run(args)` is the first method to be called on startup - Here is a simplified code snippet of its implementation:

```
public void run(String... arguments) throws Exception {

    final Bootstrap<T> bootstrap = new Bootstrap<>(this);
    bootstrap.addCommand(new ServerCommand<>(this));
    bootstrap.addCommand(new CheckCommand<>(this));

    initialize(bootstrap); // -- implemented by you; it should call:
        // 1. add bundles (typically being used)
        // 2. add commands (if any)

    // Should be called after `initialize` to give an opportunity to set a custom metric_
    ↪registry
    bootstrap.registerMetrics(); // start tracking some default jvm params...

    // for each cmd, configure parser w/ cmd
    final Cli cli = new Cli(new JarLocation(getClass()), bootstrap, our, err)
    cli.run(arguments);
}
```

Bootstrap is the pre-start (temp) application environment, containing everything required to bootstrap a Dropwizard command. Here is a simplified code snippet to illustrate its structure:

```
Bootstrap(application: Application<T>) {
    this.application = application;
    this.objectMapper = Jackson.newObjectMapper();
    this.bundles = new ArrayList<>();
    this.configuredBundles = new ArrayList<>();
    this.commands = new ArrayList<>();
    this.validatorFactory = Validators.newValidatorFactory();
}
```

(continues on next page)

(continued from previous page)

```

this.metricRegistry = new MetricRegistry();
this.classLoader = Thread.currentThread().getContextClassLoader();
this.configurationFactory = new DefaultConfigurationFactoryFactory<>();
this.healthCheckRegistry = new HealthCheckRegistry();
}

```

Environment is a longer-lived object, holding Dropwizard's Environment (not env. Such as dev or prod). It holds a similar, but somewhat different set of properties than the Bootstrap object - here is a simplified code snippet to illustrate that:

```

Environment (...) {
    // from bootstrap
    this.objectMapper = ...
    this.classLoader = ...
    this.metricRegistry = ...
    this.healthCheckRegistry = ...
    this.validator = bootstrap.getValidatorFactory().getValidator()

    // extra:
    this.bundles = new ArrayList<>();
    this.configuredBundles = new ArrayList<>();

    // sub-environments:
    this.servletEnvironment = ... // -- exposed via the servlets() method
    this.jerseyEnvironment = ... // -- exposed via the jersey() method
    this.adminEnvironment = ... // -- exposed via the admin() method
}

```

A Dropwizard `Bundle` is a reusable group of functionality (sometimes provided by the Dropwizard project itself), used to define blocks of an application's behavior. For example, `AssetBundle` from the dropwizard-assets module provides a simple way to serve static assets from your application's `src/main/resources/assets` directory as files available from `/assets/*` (or any other path) in your application.

A `ConfiguredBundle` is a bundle that requires a configuration provided by the `Configuration` object (implementing a relevant interface)

Properties such as database connection details should not be stored on the Environment; that is what your Configuration .yml file is for. Each logical environment (dev/test/staging/prod) - would have its own Configuration .yml - reflecting the differences between different "server environments".

19.1.1 Commands

Command objects are basic actions, which Dropwizard runs based on the arguments provided on the command line. The built-in `server` command, for example, spins up an HTTP server and runs your application. Each Command subclass has a name and a set of command line options which Dropwizard will use to parse the given command line arguments. The `check` command parses and validates the application's configuration.

If you will check again the first code snippet in this document - you will see creating these two commands, is the first step in the bootstrapping process.

Another important command is `db` - allowing executing various db actions, see *Dropwizard Migrations*

Similar to `ConfiguredBundle`, some commands require access to configuration parameters and should extend the `ConfiguredCommand` class, using your application's `Configuration` class as its type parameter.

19.1.2 The CLI class

Let us begin with a simplified version of the constructor:

```
public Cli(location : JarLocation, bootstrap : Bootstrap<?>,
           stdout: OutputStream, stderr: OutputStream) {
    this.stdout = stdout; this.stderr = stderr;
    this.commands = new TreeMap<>();
    this.parser = buildParser(location);
    this.bootstrap = bootstrap;
    for (command in bootstrap.commands) {
        addCommand(command)
    }
}
```

Cli is the command-line runner for Dropwizard application. Initializing, and then running it - is the last step of the Bootstrapping process.

Run would just handle commandline args (-help, -version) or runs the configured commands. E.g. - When running the server command:

```
java -jar target/hello-world-0.0.1-SNAPSHOT.jar server hello-world.yml
```

Just note the two basic commands are built of a parent, and a sub-class:

```
class CheckCommand<T extends Configuration> extends ConfiguredCommand<T>
class ServerCommand<T extends Configuration> extends EnvironmentCommand<T>
```

The order of operations is therefore:

1. Parse cmdline args, determine sub-command.
2. Run ConfiguredCommand, which get a parameter with the location of a YAML configuration file - parses and validates it.
3. CheckCommand.run() runs next, and does almost nothing: it logs "Configuration is OK"
4. Run EnvironmentCommand:
 - a) Create Environment
 - b) Calls bootstrap.run(cfg, env) - run bundles with config. & env.
 - c) Bundles run in FIFO order.
 - d) Calls application.run(cfg, env) - implemented by you
6. Now, ServerCommand.run() runs
 - a) Calls serverFactory.build(environment) - to configure Jetty and Jersey, with all relevant Dropwizard modules.
 - b) Starts Jetty.

19.2 Jetty Lifecycle

If you have a component of your app that needs to know when Jetty is going to start, you can implement `Managed` as described in the Dropwizard docs.

If you have a component that needs to be signaled that Jetty has started (this happens after all `Managed` objects' `start()` methods are called), you can register with the `env`'s lifecycle like:

```
env.lifecycle().addServerLifecycleListener(new ServerLifecycleListener() {
    @Override
    public void serverStarted(Server server) {
        // ... do things here ....
    }
});
```

CONTRIBUTORS

Dropwizard wouldn't exist without the hard work contributed by numerous individuals.

Many, many thanks to:

- _ Kilemensi
- Aaron Ingram
- Adam Jordens
- Adam Jordens
- Adam Marcus
- Adrian Suarez
- afrin216
- aharin
- Aidan
- Akhilesh Tyagi
- akumlehn
- Al Scott
- Alex Ausch
- Alex Butler
- Alex Heneveld
- Alex Katlein
- Alex Shpak
- Alexander von Renteln
- Alexey Gavrilov
- alexey-wg2
- Alice Chen
- Amr Mostafa
- Anders Hedström
- Anders Jansson
- Andreas Petersson

- Andreas Stührk
- Andrei Savu
- Andrew Clay Shafer
- AnDyXX
- anikiej
- Anna Goncharova
- Anthony Milbourne
- Anthony Wat
- Arien Kock
- Armando Singer
- Artem Grankin
- Artem Prigoda
- arunh
- Bakary DJIBA
- Bart Hanssens
- Bartek Szymański
- Basil James Whitehouse III
- Ben Bader
- Ben Blackmore
- Ben Scholl
- Ben Smith
- Benjamin Bentmann
- Benny Zlotnik
- Bo Gotthardt
- BoatMisser
- Boyd Meier
- Brandon Beck
- Brent Ryan
- Brett Hoerner
- Brian Demers
- Brian McCallister
- Brian O’Neill
- Brian Vosburgh
- Brock Mills
- Bruce Ritchie
- Bryan Burkholder

- Burak Dede
- BusComp
- Børge Nese
- Cagatay Kavukcuoglu
- Camille Fournier
- Carl Lerche
- Carlo Barbara
- Carter Kozak
- Cemalettin Koc
- Chad Selph
- Charlie Greenbacker
- Charlie La Mothe
- cheddar
- Chen W
- Chengdu Huang
- Chris Hawley
- Chris Micali
- Chris Pimlott
- Chris Tierney
- Christoffer Eide
- Christoph Kutzinski
- Christopher Cudennec
- Christopher Currie
- Christopher Currie
- Christopher Elkins
- cleiner
- Coda Hale
- Colin Armstrong
- Collin Van Dyck
- contextshuffling
- Cory Wilbur
- Craig P. Motlin
- Csaba Palfi
- dabraham02124
- Dale Wijnand
- Damian Pawlowski

- Dan Everton
- dan mcweeney
- Dan Peterström
- Dang Nguyen Anh Khoa
- Daniel
- Daniel Correia
- Daniel Temme
- Darren Yin
- David Ehrmann
- David Harris
- David Illsley
- David Martin
- David Morgantini
- David Nicholas Williams
- David Stendardi
- Denny Abraham Cheriyan
- Derek Cicerone
- Derek Stainer
- Devin Breen
- Dheerendra Rathor
- Dietrich Featherston
- Dimitris Zavalidis
- Dmitry Minkovsky
- Dmitry Ustalov
- dom farr
- Dominic Tootell
- Dominik Henneke
- Dominik Wagenknecht
- Drew Stephens
- Dylan Scott
- Eduardo Caceres
- egginbag
- Emeka Mosanya
- Emrullah YILDIRIM
- eocantu
- Erik van Oosten

- EthanLozano
- Evan Jones
- Evan Meagher
- Ezra Epstein
- Felix
- Flemming Frandsen
- Florian Hirsch
- florinn
- Francisco Rojas
- Fred Deschenes
- Fredrik Sundberg
- Friso Vrolijk
- Frode Nerbråten
- Fábio Franco Uechi
- Gabe Henkes
- Gary Dusbabek
- Gautam Gupta
- Glenn McAllister
- Graham O'Regan
- Grzegorz Rożniecki
- Guillaume Simard
- Gunnar Ahlberg
- Harry Howarth
- Henrik Hegardt
- Henrik S.
- Hernan Silberman
- Hrvoje Slaviček
- Hugo Gonçalves
- Håkan Jonson
- Ian Eure
- Ian Ferguson
- Ian White
- Igor Savin
- Ilias Bartolini
- islasjuanp
- Jacek Jackowiak

- Jake Swenson
- James Morris
- James Alfei
- James Ward
- Jamie Furnaghan
- Jan Galinski
- Jan-Terje Sørensen
- Jared Stehler
- Jason
- Jason Clawson
- Jason Dunkelberger
- Jason Toffaletti
- Jasper Vandemalle
- Javier Campanini
- jduan
- Jeff Klukas
- Jelmer ter Wal
- Jerry-Carter
- Jilles Oldenbeuving
- Jochen Schalanda
- Joe Barnett
- Joe Lauer
- Joe Schmetzer
- Johan Wirde (@jwirde)
- Jon Radon
- Jonathan Haber
- Jonathan Halterman
- Jonathan Monette
- Jonathan Ruckwood
- Jonathan Welzel
- Jordan Moore
- Jordan Zimmerman
- Joshua Spiewak
- Julien
- Justin Miller
- Justin Plock

- Jérémie Panzer
- Kashyap Paidimarri
- Kirill Vlasov
- Konstantin Yegupov
- Koray Tugay
- Kristian Klette
- Kristian Schjelderup
- Krzysztof Mejka
- LeekAnarchism
- Leo Fernandes
- leoconco
- Lior Bar-On
- Lucas Pleß
- Lunfu Zhong
- maffe
- Magnus Reftel
- Maher Abuthraa
- Malte S. Stretz
- Manabu Matsuzaki
- Manuel Hegner
- Marcin Biegan
- Marcos Paulo Belasco de Almeida
- Marcus Höjvall
- Marius Volkhart
- Mark Elliot
- Mark Lilback
- Mark Reddy
- Mark Symons
- Mark Wolfe
- markez92
- Martin W. Kirst
- Matt Brown
- Matt Carrier
- Matt Hurne
- Matt Nelson
- Matt Snider

- Matt Veitas
- Matt Whipple
- Matthew Clarke
- Matthew Dolan
- Matthew Simoneau
- Matthias Bläsing
- Matthias Müller
- Max Wenzin
- Maximilien Marie
- Michael Chaten
- Michael Fairley
- Michael McCarthy
- Michael Rice
- Michael Silvanovich
- Michal Rutkowski
- Michel Decima
- Michiel Leegwater
- Miguel Ángel Martín
- Mikael Amborn
- Mike Miller
- Mikhail Gromov
- Mohammad Nasir Rasul
- Moritz Kammerer
- Moxie Marlinspike
- Mårten Gustafson
- nanreh
- Nasir
- Natalie Zamani
- Natan Abolafya
- Nicholas Heitz
- Nick Babcock
- Nick Smith
- Nick Telford
- Nikhil Bafna
- Nisarg Shah
- Oddmar Sandvik

- Oleg
- Oliver B. Fischer
- Olivier Abdesselam
- Olivier Chédru
- Olivier Grégoire
- Ori Schwartz
- Oscar Nalin Nilsson
- Otto Jongerius
- Owen Jacobson
- pandaadb
- Pablo Fernandez
- Patrick Stegmann
- Patryk Najda
- Paul Kenneth Kent
- Paul Samsatha
- Paul Tomlin
- Peter Sear
- Peter Stackle
- Peter Wippermann
- Philip K. Warren
- Philip Potter
- pkokush
- Punyashloka Biswal
- Qinfeng Chen
- Quoc-Viet Nguyen
- Rachel Normand
- Radoslav Petrov
- RawToast
- Richard Kettelerij
- Richard Nyström
- Robert Barbey
- Rohan Nagar
- Russell Howe
- Ryan Berdeen
- Ryan Kennedy
- Ryan Warren

- Rémi Alvergnat
- Rüdiger zu Dohna
- Sam Quigley
- Satoshi “Moris” Tagomori
- Scott D.
- Sean Scanlon
- Sergio Escalante
- shail
- Sharath B. Patel
- shartte
- Shawn Smith
- Simon Collins
- Simon Dean
- Sjoerd Talsma
- smolloy
- Sourav Mitra
- Stan Svec
- Stephen Huenneke
- Stephen Souness
- Steve Agalloco
- Steve Hill
- Steven Benitez
- Stevo Slavić
- Stuart Gunter
- sullis
- Suryatej Mukkamalla
- Szymon Pacanowski
- Takuma Tsutsumi
- Tatu Saloranta
- Tatu Saloranta
- Ted Nyman
- terezivy
- Thiago Moretto
- Thomas Sundberg
- Tim Bart
- Tim Bartley

- Timothée Peignier
- Todd Bednarczyk
- Tom Akehurst
- Tom Crayford
- Tom Lee
- Tom Morris
- Tom Shen
- Tomas Crhak
- Tomasz Adamski
- Tony Gaetani
- Trevor Mack
- Tristan Buckner
- Tristan Burch
- twilson-palantir
- Tyrone Cutajar
- Vadim Spivak
- vanvlack
- Varun Loiwal
- Vasyl Vavrychuk
- Victor Noël
- Vitor Reis
- Vladimir Ladynev
- Vojtěch Vondra
- vzx
- Will Sommers
- William Herbert
- William Palmer
- Xavier Shay
- Xiaodong Xie
- Yaroslav Admin
- YE Qing
- Yiwei Gao
- Yun Zhi Lin
- Yurii Savka
- zebra-kangaroo
- zmarois

SPONSORS

Dropwizard is generously supported by some companies with licenses and free accounts for their products.

21.1 JetBrains



JetBrains supports our open source project by sponsoring some [All Products Packs](#) within their [Free Open Source License](#) program.

OTHER VERSIONS

- 1.3.x
- 1.2.x
- 1.1.x
- 1.0.x
- 0.9.x
- 0.8.x
- 0.7.x
- 0.6.2