



WHITEPAPER

Why and how should you optimize Docker images with Jib?

As enterprises across industries look to fast-track their digital transformation journeys, business leaders are increasingly looking to reduce their application startup time and reduce general costs over unwanted network and internet usage. Towards that end, software developers and DevOps engineers are now looking to build, run, and share applications with containers. However, as they do so, they are running into some issues.

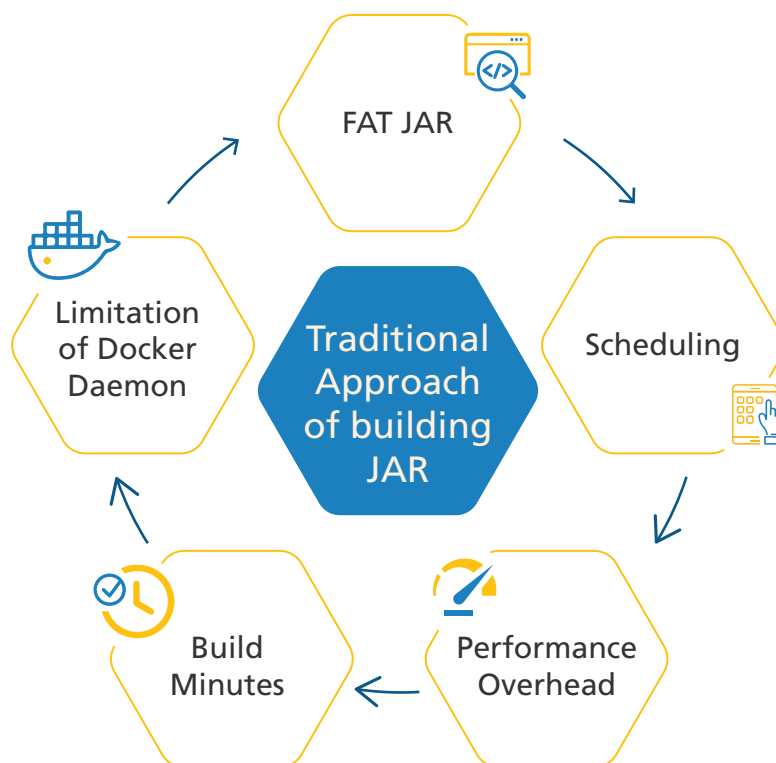


The key challenges for Java developers, storage and long build times

Currently, most Java based applications are using Spring Boot in Docker with Maven and are creating standard fat Boot jars with all the dependencies packed inside. In this process, a new 250 MB image is created each time anything is built, even if it's very little code change, resulting in **inefficient storage space usage in our private repository**. This is because the fat jar contains both shared dependencies (which change infrequently) and our code.

In a container orchestration system, the container image is pulled from the image registry to a host running a container engine. This process is called scheduling. Pulling large-sized images from the registry result in long scheduling times in container orchestration systems and **long build times** in CI pipelines.

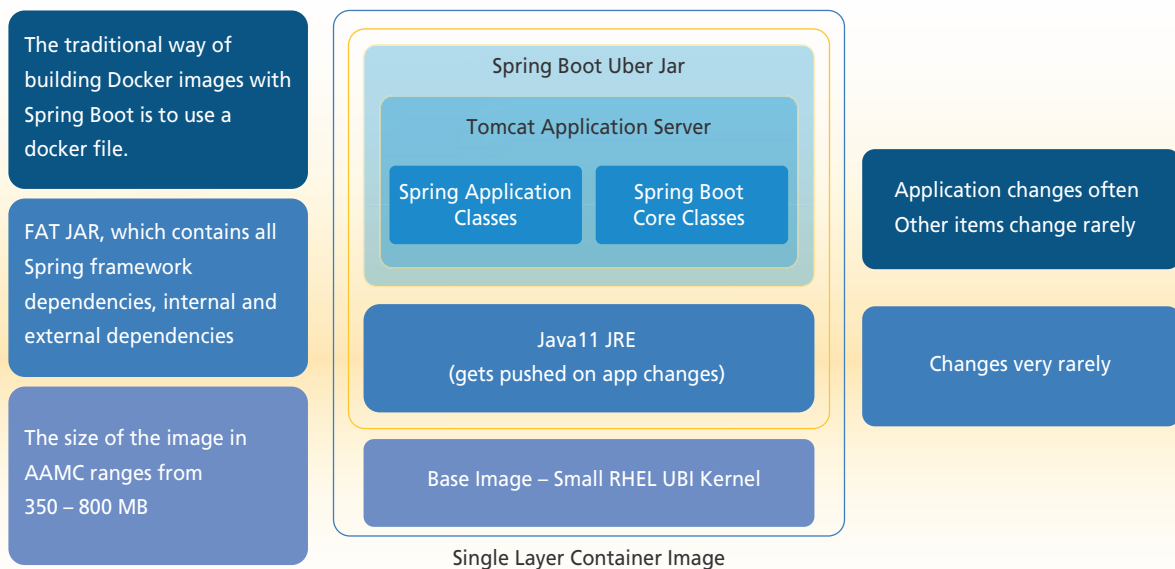
Large docker images and frequent release cycles **can eat up a lot of gigabytes of storage space and increase network traffic**. This can be potentially poor design, especially for cloud designs where you pay for traffic usage and storage. Pulling 1GB of data from your registry every single time results in long scheduling processes and long build times in CI pipelines.



When you multiply this scenario in context to all our pipelines, think about the **time the CI systems will need** to place this container and then multiply with the number of applications you have. Further, you have an additional step of first creating a docker image and then pushing it to Artifactory/ECR, again resulting in the consumption of increased build minutes.

Most organizations also have a **problem installing docker on the developer workspace**. Hence testing these builds will need to be deployed into systems that have docker installed and then validated, which can be time-consuming during critical times.

Traditional Builds



So, what's the solution?

We can resolve the above problem using two different approaches.

- Building and deploying only the delta
- Abstract image creation

1. Building and deploying only the delta

The process of deploying changes to an existing application is called [delta deployment](#). Delta deployment supports several scenarios in which a full deployment process would be inefficient. Typically, we deploy the parts of the application that are frequently updated.

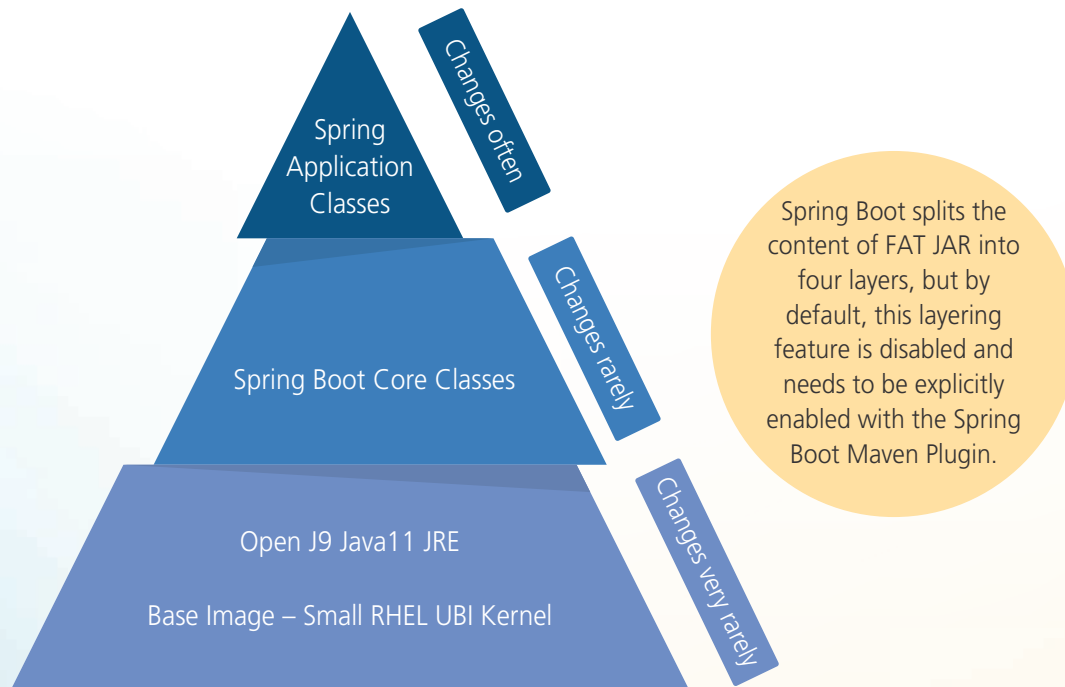
Using Layering Feature in Spring

Spring Boot splits the content of *fat Jar* into 4 layers. However, by default, this layering feature is disabled and needs to be explicitly enabled with the Spring Boot Maven plugin.

Maven Plugin

```
<plugin>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-maven-plugin</artifactId>
  <configuration>
    <layers>
      <enabled>true</enabled>
    </layers>
  </configuration>
</plugin>
```

Layering to the Solution



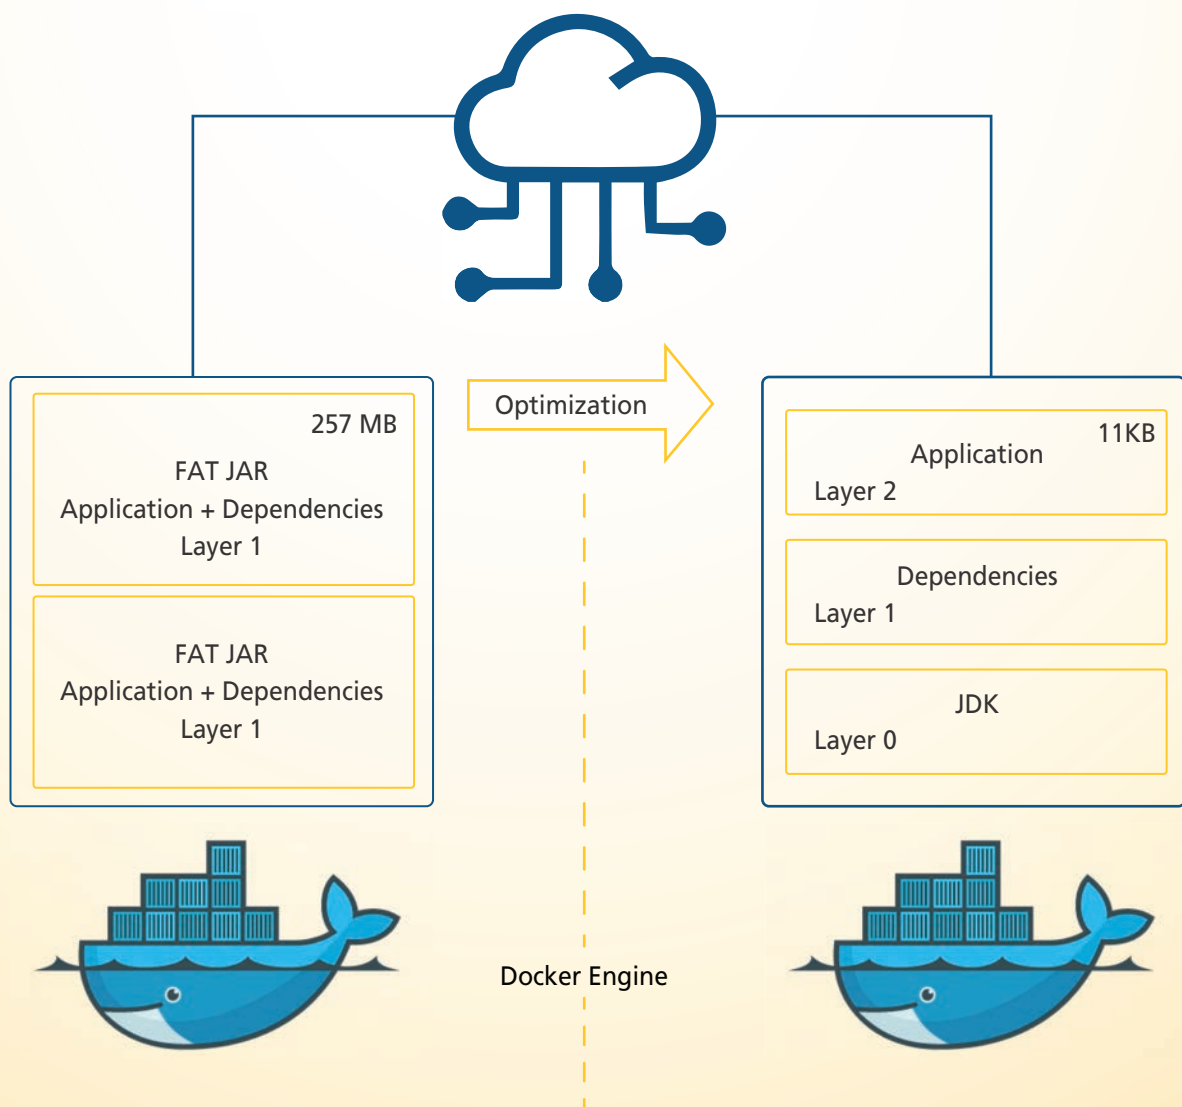
The default layers are:

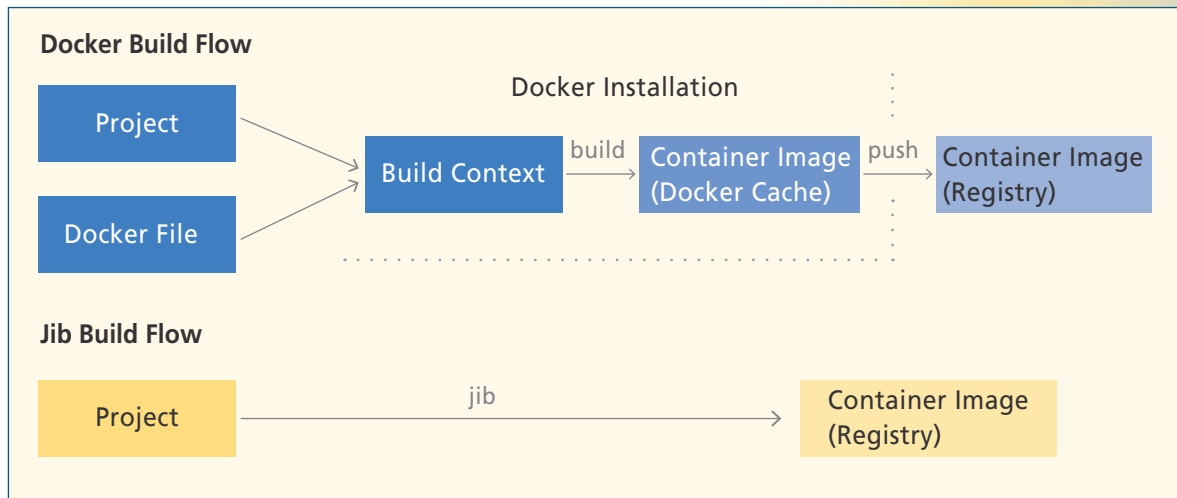
Layer Name	Contents
Application	Application classes and resources
Snapshot Dependencies	Any dependency, whose version contains SNAPSHOT
Spring-Boot-Loader	JAR loader classes
Dependencies	Any dependency, whose version does not contain SNAPSHOT

The layers are defined in a `layers.idx` file in the order that they should be added to the Docker image. These layers get cached in the host after the first pull since they do not change. Only the updated application layer is downloaded to the host, which is faster because of the reduced size.

Building the Image with Dependencies Extracted in Separate Layers

We will build the final image in two stages using a method called multi-stage build. In the first stage, we will extract the dependencies, and in the second stage, we will copy the extracted dependencies to the final image.





Jib is available as a plugin for Maven

(<https://github.com/GoogleContainerTools/jib/tree/master/jib-maven-plugin>) and Gradle

(<https://github.com/GoogleContainerTools/jib/tree/master/jib-gradle-plugin>) and requires minimal

configuration. Simply add the plugin to your build definition and configure the target image. If you are

building to a private registry, make sure to [configure Jib with credentials for your registry.](#)

(<https://github.com/GoogleContainerTools/jib/tree/master/jib-maven-plugin#authentication-methods>) The

easiest way to do this is to use credential helpers like [docker-credential-gcr](#)

(<https://github.com/GoogleContainerTools/jib/blob/master/docs/configure-gcp-credentials.md>). Jib also

provides additional rules for building an image to a Docker daemon if you need it.

Use Case

For this example, we'll provide our Docker Hub credentials to `.m2/settings.xml`:

```

1 <servers>
2   <server>
3     <id>registry.hub.docker.com</id>
4     <username><DockerHub Username></username>
5     <password><DockerHub Password></password>
6   </server>
7 </servers>

```

Deploying to Jfrog with Jib:

Now, we can use `jib-maven-plugin`, or the Gradle equivalent, to containerize our application with a simple command:

```
mvn compile com.google.cloud.tools:jib-maven-plugin:2.5.0:build -Dimage=$IMAGE_PATH
```

where `IMAGE_PATH` is the target path in the container registry.

For example, to upload the image `spring-jib-app` to Jfrog, we would do:

```
export IMAGE_PATH=registry.hub.docker.com/spring-jib-app
```

And that's it! This will build the docker image of our application and push it to the Jfrog.

We can, of course, upload the image to [Google Container Registry](https://cloud.google.com/container-registry/) (<https://cloud.google.com/container-registry/>) or [Amazon Elastic Container Registry](https://aws.amazon.com/ecr/) (<https://aws.amazon.com/ecr/>) in a similar way.

Simplifying the Maven Command

Also, we can shorten our initial command by configuring the plugin in our pom instead, like any other Maven plugin.

```
1 <project>
2   ...
3   <build>
4     <plugins>
5       ...
6       <plugin>
7         <groupId>com.google.cloud.tools</groupId>
8         <artifactId>jib-maven-plugin</artifactId>
9         <version>2.5.0</version>
10        <configuration>
11          <to>
12            <image>${image.path}</image>
13          </to>
14        </configuration>
15      </plugin>
16      ...
17    </plugins>
18  </build>
19  ...
20 </project>
```


With this change, we can simplify our maven command:

Now, we can use jib-maven-plugin, or the Gradle equivalent, to containerize our application with a simple command:

```
mvn compile jib:build
```

When we build the application, an image is created, and it is stored in Jfrog.

Benefits of this approach:

- Optimization will **help teams share smaller images**, improve performance, and make it easier to debug problems.
- Docker caches images. If you need to create a custom base image, layers optimization, and multiple instances of the same layers, it will **speed up the load times and make it easier to track**.
- **Bypassing Docker daemon** to be installed on developer machines is a boon for companies that have restrictions to install docker on developer machines.
- JIB layers application binaries and thus deploys only layers only that change, thus **reducing network bandwidth** and cost it incurred before to download the entire fat jar and deploy.
- All of the image creation and downloading are usually done via the CI platform. This further increases the build minutes. Hence since JIB takes over and does this for us, we are **effectively reducing the build minutes**, which makes this approach **cost-effective**.

In a nutshell:

Effective use of Jib containerizing applications has proven to reduce the scheduling time of the container orchestration system. This has further resulted in reducing the application start-up time. Further reducing the application start-up time. This approach also removes manual docker file creation, which is not straightforward and can bring in adversities in application security. It also optimizes the containers by enabling the layering feature, which extracts the dependencies in separate layers that get cached in the host, and the thin layer of the application is downloaded during scheduling in container runtime engines.



Santosh Malimath

Sr Project Manager/Technical Architect

Over the last 16 years, Santosh has worked with several large enterprises. His experience spans as a Full-stack development, DevSecOps engineering and IOT solutioning. During the last 6 years he has engaged in analyzing and creating project scope and milestones for several technical company initiatives. Has been contributing to help define key partnerships to targets, establishing technical relationships, and managing the day-to-day interactions to build long-term business and marketing opportunities.

About LTIMindtree

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700+ clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by nearly 90,000 talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit www.ltimindtree.com.