

# Java Persistence API (JPA)

jim stafford

Fall 2024 v2022-07-24: Built: 2024-11-19 21:34 EST

# Table of Contents

1. Introduction	1
1.1. Goals	1
1.2. Objectives	1
2. Java Persistence API	2
2.1. JPA Standard and Providers	2
2.2. Javax / Jakarta Transition	2
2.3. JPA Dependencies	3
2.4. Enabling JPA AutoConfiguration	4
2.5. Configuring JPA DataSource	5
2.6. Automatic Schema Generation	5
2.7. Schema Generation to File	5
2.8. Generated Schema Output	6
2.9. Other Useful Properties	7
2.10. Configuring JPA Entity Scan	8
2.11. JPA Persistence Unit	9
2.12. JPA Persistence Context	9
3. JPA Entity	11
3.1. JPA @Entity Defaults	11
3.2. JPA Overrides	12
4. Basic JPA CRUD Commands	14
4.1. EntityManager persist()	14
4.2. EntityManager find() By Identity	15
4.3. EntityManager query	15
4.4. EntityManager flush()	16
4.5. EntityManager remove()	17
4.6. EntityManager clear() and detach()	17
5. Transactions	19
5.1. Transactions Required for Explicit Changes/Actions	19
5.2. Activating Transactions	19
5.3. Conceptual Transaction Handling	20
5.4. Activating Transactions in @Components	21
5.5. Calling @Transactional @Component Methods	21
5.6. @Transactional @Component Methods SQL	22
5.7. Unmanaged @Entity	22
5.8. Shared Transaction	23
5.9. @Transactional Attributes	23
6. Summary	25

# Chapter 1. Introduction

This lecture covers implementing object/relational mapping (ORM) to an RDBMS using the Java Persistence API (JPA). This lecture will directly build on the previous concepts covered in the RDBMS and show the productivity power gained by using an ORM to map Java classes to the database.

## 1.1. Goals

The student will learn:

- to identify the underlying JPA constructs that are the basis of Spring Data JPA Repositories
- to implement a JPA application with basic CRUD capabilities
- to understand the significance of transactions when interacting with JPA

## 1.2. Objectives

At the conclusion of this lecture and related exercises, the student will be able to:

1. declare project dependencies required for using JPA
2. define a `DataSource` to interface with the RDBMS
3. define a Persistence Context containing an `@Entity` class
4. inject an `EntityManager` to perform actions on a PersistenceUnit and database
5. map a simple `@Entity` class to the database using JPA mapping annotations
6. perform basic database CRUD operations on an `@Entity`
7. define transaction scopes

# Chapter 2. Java Persistence API

The Java Persistence API (JPA) is an object/relational mapping (ORM) layer that sits between the application code and JDBC and is the basis for Spring Data JPA Repositories. JPA permits the application to primarily interact with plain old Java (POJO) business objects and a few standard persistence interfaces from JPA to fully manage our objects in the database. JPA works off convention and customized by annotations primarily on the POJO, called an Entity. JPA offers a rich set of capabilities that would take us many chapters and weeks to cover. I will just cover the very basic setup and `@Entity` mapping at this point.

## 2.1. JPA Standard and Providers

The JPA standard was originally part of Java EE, which is now managed by the [Eclipse Foundation within Jakarta](#). It was released just after Java 5, which was the first version of Java to support annotations. It replaced the older, heavyweight Entity Bean Standard — that was ill-suited for the job of realistic O/R mapping — and progressed on a path in line with Hibernate. There are several [persistence providers of the API](#)

- [EclipseLink](#) is now the reference implementation
- [Hibernate](#) was one of the original implementations and the default implementation within Spring Boot
- [DataNucleus](#)

## 2.2. Javax / Jakarta Transition

JPA transitioned from Oracle/JavaEE to Jakarta in ~2019 and was officially renamed from "Java Persistence API" to "Jakarta Persistence". Official references dropped the "API" portion of the name and used the long name. Unofficial references keep the "API" portion of the name and still reference the product as "JPA". For simplicity — I will take the unofficial reference route and continue to refer to the product as "JPA" for short once we finish this background paragraph.

The [last release](#) of the "Java Persistence API" (officially known as "JPA") was 2.2 in ~2017. You will never see a version of the "Java Persistence API" newer than 2.2.x. "Jakarta Persistence" (unofficially known as "JPA") [released a clone](#) of the "Java Persistence API" version 2.2.x under the [jakarta](#) Maven package naming to help in the transition. "Jakarta Persistence" 3.0 was released in ~2020 with just package renaming from `javax.persistence` to `jakarta.persistence`. Enhancements were not added until JPA 3.1 in ~2022. That means the feature set remained idle for close to 5 years during the transition.

### *JPA Maven Module and Java Packaging Transition*



- `javax.persistence:javax.persistence-api`
  - used `javax.persistence` Java package naming
  - ended with version 2.2 in 2017
- `jakarta.persistence:jakarta.persistence-api 2.2.x`

- used `javax.persistence` Java package naming
- released in 2018
- `jakarta.persistence:jakarta.persistence-api 3.x`
  - uses `jakarta.persistence` Java package naming
  - released in 2020

Spring Boot 2.7 advanced to `jakarta.persistence:jakarta.persistence-api:jar:2.2.3`, which kept the `javax.persistence` Java package naming. Spring Boot 3 advanced to `jakarta.persistence:jakarta.persistence-api:jar:3.x.x`, which required all Java package references to change from `javax.persistence` to `jakarta.persistence`. With a port to Spring Boot 3/Spring 6, we have finally turned the corner on the `javax/jakarta` naming issues.

#### *Spring Boot Javax/Jakarta Transition*



- Spring Boot 2.7
  - uses `jakarta.persistence:jakarta.persistence-api 2.2.x`
  - with `javax.persistence` Java package naming
- Spring Boot 3
  - uses `jakarta.persistence:jakarta.persistence-api 3.x`
  - with `jakarta.persistence` Java package naming

## 2.3. JPA Dependencies

Access to JPA requires declaring a dependency on the JPA interface (`jakarta.persistence-api`) and a provider implementation (e.g., `hibernate-core`). This is automatically added to the project by declaring a dependency on the `spring-boot-starter-data-jpa` module.

#### *Spring Data JPA Maven Dependency*

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

The following shows a subset of the dependencies brought into the application by declaring a dependency on the JPA starter.

#### *Spring Boot Starter JPA 3.1 Dependencies*

```
+-- org.springframework.boot:spring-boot-starter-data-jpa:jar:3.3.2:compile
|  +- org.springframework.boot:spring-boot-starter-jdbc:jar:3.3.2:compile
|  ...
|  | \- org.springframework:spring-jdbc:jar:6.1.11:compile
|  +- org.hibernate.orm:hibernate-core:jar:6.5.2.Final:compile ②
|  | +- jakarta.persistence:jakarta.persistence-api:jar:3.1.0:compile ①
```

```
| | +- jakarta.transaction:jakarta.transaction-api:jar:2.0.1:compile
...

```

- ① the JPA API module is required to compile standard JPA constructs
- ② a JPA provider module is required to access extensions and for runtime implementation of the standard JPA constructs

#### Spring Boot Starter JPA 2.7 Dependencies

```
+ org.springframework.boot:spring-boot-starter-data-jpa:jar:2.7.0:compile
| +- org.springframework.boot:spring-boot-starter-jdbc:jar:2.7.0:compile
| | \- org.springframework:spring-jdbc:jar:5.3.20:compile
| +- jakarta.persistence:jakarta.persistence-api:jar:2.2.3:compile ①
| +- org.hibernate:hibernate-core:jar:5.6.9.Final:compile ②
...

```

- ① the JPA API module is required to compile standard JPA constructs
- ② a JPA provider module is required to access extensions and for runtime implementation of the standard JPA constructs

From these dependencies, we can define and inject various JPA beans.

## 2.4. Enabling JPA AutoConfiguration

JPA has its own defined bootstrapping constructs that involve settings in `persistence.xml` and entity mappings in `orm.xml` configuration files. These files define the overall persistence unit and include information to connect to the database and any custom entity mapping overrides.

Spring Boot JPA automatically configures a default persistence unit and other related beans when the `@EnableJpaRepositories` annotation is provided. `@EntityScan` is used to identify packages for `@Entities` to include in the persistence unit.

#### Spring Boot Data Bootstrapping

```
import org.springframework.data.jpa.repository.config.EnableJpaRepositories;

@SpringBootApplication
@EnableJpaRepositories ①
// Class<?>[] basePackageClasses() default {};
// String repositoryImplementationPostfix() default "Impl";
// ...(many more configurations)
@EntityScan ②
// Class<?>[] basePackageClasses() default {};
public class JPASongsApp {

```

- ① triggers and configures scanning for JPA Repositories (the topic of the next lecture)
- ② triggers and configures scanning for JPA Entities

By default, this configuration will scan packages below the class annotated with the `@EntityScan` annotation. We can override that default using the attributes of the `@EntityScan` annotation.

## 2.5. Configuring JPA DataSource

Spring Boot provides convenient ways to provide property-based configurations through its standard property handling, making the connection areas of `persistence.xml` unnecessary (but still usable). The following examples show how our definition of the `DataSource` for the JDBC/SQL example can be used for JPA as well.

Table 1. Spring Data JPA Database Connection Properties

### H2 In-Memory Example Properties

```
spring.datasource.url=jdbc:h2:mem:songs
```

### Postgres Client Example Properties

```
spring.datasource.url=jdbc:postgresql://localhost:5432/postgres
spring.datasource.username=postgres
spring.datasource.password=secret
```

## 2.6. Automatic Schema Generation

JPA provides the capability to automatically generate schema from the Persistence Unit definitions. This can be configured to write to a file to be used to kickstart schema authoring. However, the most convenient use for schema generation is at runtime during development.

Spring Boot will automatically enable runtime schema generation for in-memory database URLs. We can also explicitly enable runtime schema generation using the following Hibernate property.

### Example Explicit Enable Runtime Schema Generation

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

## 2.7. Schema Generation to File

The JPA provider can be configured to generate schema to a file. This can be used directly by tools like Flyway or simply to kickstart manual schema authoring.

The following configuration snippet instructs the JPA provider to generate a create and drop commands into the same `drop_create.sql` file based on the metadata discovered within the Persistence Context. Hibernate has the additional features to allow for formatting and line termination specification.

### Schema Generation to File Example

```
spring.jpa.properties.jakarta.persistence.schema-generation.scripts.action=drop-and-
```

```

create
spring.jpa.properties.jakarta.persistence.schema-generation.create-source=metadata

spring.jpa.properties.jakarta.persistence.schema-generation.scripts.create-
target=target/generated-sources/ddl/drop_create.sql
spring.jpa.properties.jakarta.persistence.schema-generation.scripts.drop-
target=target/generated-sources/ddl/drop_create.sql

spring.jpa.properties.hibernate.hbm2ddl.delimiter=; ①
spring.jpa.properties.hibernate.format_sql=true ②

```

① adds ";" character to terminate every command — making it SQL script-ready

② adds new lines to make more human-readable

`action` can have values of `none`, `create`, `drop-and-create`, and `drop` <sup>[1]</sup>

`create/drop-source` can have values of `metadata`, `script`, `metadata-then-script`, or `script-then-metadata`. `metadata` will come from the class defaults and annotations. `script` will come from a location referenced by `create/drop-script-source`



#### *Generate Schema to Debug Complex Mappings*

Generating schema from `@Entity` class metadata is a good way to debug odd persistence behavior. Even if normally ignored, the generated schema can identify incorrect and accidental definitions that may cause unwanted behavior. I highly recommend activating it while working with new `@Entity` mappings or debugging old ones.

## 2.8. Generated Schema Output

The following snippet shows the output of the Entity Manager when provided with the property activations and configuration. The specific output contains all SQL commands required to drop existing rows/schema (by the same name(s)) and instantiate a new version based upon the `@Entity` metadata.

### *Generated Schema Output*

```

$ docker compose up postgres -d ①
$ mvn clean package spring-boot:start spring-boot:stop \
-Dspring-boot.run.profiles=postgres -DskipTests ②

$ cat target/generated-sources/ddl/drop_create.sql ③

drop table if exists reposongs_song cascade;

drop sequence if exists reposongs_song_sequence;

create sequence reposongs_song_sequence start with 1 increment by 50;

create table reposongs_song (

```



```
id integer not null,  
released date,  
artist varchar(255),  
title varchar(255),  
primary key (id)  
);
```

- ① starting Postgres DB
- ② starting/stopping server with JPA EntityManager
- ③ JPA EntityManager-produced schema file

## 2.9. Other Useful Properties

It is useful to see database SQL commands coming from the JPA/Hibernate layer during early stages of development or learning. The following properties will print the JPA SQL commands and values mapped to the SQL substitution variables.

The first two property settings both functionally produce logging of SQL statements.

*Two JPA/Hibernate SQL Logging Options*

```
spring.jpa.show-sql=true  
logging.level.org.hibernate.SQL=DEBUG
```

- The `spring.jpa.show-sql` property controls raw text output, lacking any extra clutter. This is a good choice if you want to see the SQL command and do not need a timestamp or threadId.

*spring.jpa.show-sql=true*

```
Hibernate: delete from REPOSONGS_SONG  
Hibernate: select next value for reposongs_song_sequence  
Hibernate: insert into reposongs_song (artist,released,title,id) values (?, ?, ?, ?)  
Hibernate: select count(s1_0.id) from reposongs_song s1_0 where s1_0.id=?
```

- The `logging.level.org.hibernate.SQL` property controls SQL logging passed through the logging framework. This is a good choice if you are looking at a busy log with many concurrent threads. The threadId provides you with the ability to filter the associated SQL commands. The timestamps provide a basic ability to judge performance timing.

*logging.level.org.hibernate.SQL=debug*

```
09:53:38.632 main DEBUG org.hibernate.SQL#logStatement:135 delete from  
REPOSONGS_SONG  
09:53:38.720 main DEBUG org.hibernate.SQL#logStatement:135 select next value for  
reposongs_song_sequence  
09:53:38.742 main DEBUG org.hibernate.SQL#logStatement:135 insert into  
reposongs_song (artist,released,title,id) values (?, ?, ?, ?)  
09:53:38.790 main DEBUG org.hibernate.SQL#logStatement:135 select count(s1_0.id)
```

```
from reposongs_song s1_0 where s1_0.id=?
```

An additional property can be defined to show the mapping of data to/from the database rows. The `logging.level.org.hibernate.orm.jdbc.bind` property controls whether Hibernate prints the value of binding parameters (e.g., "?") input to or returning from SQL commands. It is quite helpful if you are not getting the field data you expect.

#### JPA/Hibernate JDBC Parameter Mapping Debug Property

```
logging.level.org.hibernate.orm.jdbc.bind=TRACE
```

The following snippet shows the result information of the activated debug. We can see the individual SQL commands issued to the database as well as the parameter values used in the call and extracted from the response. The extra logging properties have been redacted from the output.

`logging.level.org.hibernate.orm.jdbc.bind=trace`

```
insert into reposongs_song (artist,released,title,id) values (?,?,,?)

binding parameter (1:VARCHAR) <- [The Orb]
binding parameter (2:DATE) <- [2020-08-01]
binding parameter (3:VARCHAR) <- [Mother Night teal]
binding parameter (4:INTEGER) <- [1]
```



#### Bind Logging Moved Packages

Hibernate moved some classes. `org.hibernate.type` logging changed to `org.hibernate.orm.jdbc.bind` in version 6.

```
logging.level.org.hibernate.type=TRACE           #older
logging.level.org.hibernate.orm.jdbc.bind=TRACE #newer
```

## 2.10. Configuring JPA Entity Scan

Spring Boot JPA will automatically scan for `@Entity` classes. We can provide a specification to external packages to scan using the `@EntityScan` annotation.

The following shows an example of using a String package specification to a root package to scan for `@Entity` classes.

#### @EntityScan example

```
import org.springframework.boot.autoconfigure.domain.EntityScan;
...
@EntityScan(value={"info.ejava.examples.db.repo.jpa.songs.bo"})
```

The following example, instead uses a Java class to express a package to scan. We are using a

specific `@Entity` class in this case, but some may define an interface simply to help mark the package and use that instead. The advantage of using a Java class/interface is that it will work better when refactoring.

*@EntityScan .class Example*

```
import info.ejava.examples.db.repo.jpa.songs.bo.Song;
...
@EntityScan(basePackageClasses = {Song.class})
```

## 2.11. JPA Persistence Unit

The JPA Persistence Unit represents the overall definition of a group of Entities and how we interact with the database. A defined Persistence Unit can be injected into the application using an `EntityManagerFactory`. From this injected instance, clients can gain access to metadata and initiate a Persistence Context.

*Persistence Unit/EntityManagerFactory Injection Example*

```
import jakarta.persistence.EntityManagerFactory;
...
@Autowired
private EntityManagerFactory emf;
```

It is rare that you will ever need an `EntityManagerFactory` and could be a sign that you may not understand yet what it is meant for versus an injected `EntityManager`. `EntityManagerFactory` is used to create a custom persistence context and transaction. The caller is responsible for beginning and committing the transaction—just like with JDBC. This sometimes is useful to create a well-scoped transaction within a transaction. It is a manual process, outside the management of Spring.

Your first choice should be to inject an `EntityManager`.



*Primarily use EntityManager versus EntityManagerFactory*

It is rare that one requires the use of an `EntityManagerFactory` and using one places many responsibilities on the caller "to do the right thing". This should be used for custom transactions. Normal business transactions should be handled using a JPA Persistence Context, with the injection of an `EntityManager`.

## 2.12. JPA Persistence Context

A Persistence Context is a usage instance of a Persistence Unit and is represented by an `EntityManager`. An injected `EntityManager` provides a first-layer cache of entities for business logic to share and transaction(s) to commit or rollback their state together. An `@Entity` with the same identity is represented by a single instance within a Persistence Context/`EntityManager`.

## Persistence Context/EntityManager Injection Example

```
import jakarta.persistence.EntityManager;  
...  
@Autowired  
private EntityManager em;
```

Injected `EntityManager`s reference the same Persistence Context when called within the same thread. That means that a `Song` loaded by one client with ID=1 will be available to sibling code when using ID=1.



### *Use/Inject EntityManager*

Normal application code that creates, gets, updates, and deletes `@Entity` data should use an injected `EntityManager` and allow the transaction management to occur at a higher level.

[1] "JavaEE: The JavaEE Tutorial, Database Schema Creation", Oracle, JavaEE 7

# Chapter 3. JPA Entity

A JPA `@Entity` is a class mapped to the database that primarily represents a row in a table. The following snippet is the example `Song` class we have already manually mapped to the `REPOSONGS_SONG` database table using manually written schema and JDBC/SQL commands in a previous lecture. To make the class an `@Entity`, we must:

- annotate the class with `@Entity`
- provide a no-argument constructor
- identify one or more columns to represent the primary key using the `@Id` annotation
- override any convention defaults with further annotations

## JPA Example Entity

```
@jakarta.persistence.Entity ①
@Getter
@AllArgsConstructor
@NoArgsConstructor ②
public class Song {
    @jakarta.persistence.Id ③ ④
    private int id;
    @Setter
    private String title;
    @Setter
    private String artist;
    @Setter
    private java.time.LocalDate released;
}
```

- ① class must be annotated with `@Entity`
- ② class must have a no-argument constructor
- ③ class must have one or more fields designated as the primary key
- ④ annotations can be on the field or property and the choice for `@Id` determines the default



### *Primary Key property is not modifiable*

This Java class is not providing a setter for the field mapped to the primary key in the database. The primary key will be generated by the persistence provider at runtime and assigned to the field. The field cannot be modified while the instance is managed by the provider. The all-args constructor can be used to instantiate a new object with a specific primary key.

## 3.1. JPA `@Entity` Defaults

By convention and supplied annotations, the class as shown above would:

- have the entity name "Song" (important when expressing JPA queries; ex. `select s from Song s`)

- be mapped to the `SONG` table to match the entity name
- have columns `id integer`, `title varchar`, `artist varchar`, and `released (date)`
- use `id` as its primary key and manage that using a provider-default mechanism

## 3.2. JPA Overrides

All the convention defaults can be customized by further annotations. We commonly need to:

- supply a table name that matches our intended schema (i.e., `select * from REPOSONGS_SONG` vs `select * from SONG`)
- select which primary key mechanism is appropriate for our use. JPA 3 no longer allows an unnamed generator for a `SEQUENCE`.
- define the `SEQUENCE` to be consistent with our SQL definition earlier
- supply column names that match our intended schema
- identify which properties are optional, part of the initial `INSERT`, and `UPDATE` -able
- supply other parameters useful for schema generation (e.g., String length)

*Common JPA Annotation Overrides*

```

@Entity
@Table(name="REPOSONGS_SONG") ①
@NoArgsConstructor
...
@SequenceGenerator(name="REPOSONGS_SONG_SEQUENCE", allocationSize = 50)③
public class Song {
    @Id
    @GeneratedValue(strategy = GenerationType.SEQUENCE,
                    generator = "REPOSONGS_SONG_SEQUENCE")②
    @Column(name = "ID") ④
    private int id;
    @Column(name="TITLE", length=255, nullable=true, insertable=true, updatable=true)⑤
    private String title;
    private String artist;
    private LocalDate released;
}

```

- ① overriding the default table name `SONG` with `REPOSONGS_SONG`
- ② overriding the default primary key mechanism with `SEQUENCE`. JPA 3 no longer permits an unnamed sequence generator.
- ③ defining the `SEQUENCE`
- ④ re-asserting the default convention column name `ID` for the `id` field
- ⑤ re-asserting many of the default convention column mappings



*Schema generation properties aren't used at runtime*

Properties like `length` and `nullable` are only used during optional JPA schema generation and are not used at runtime.

# Chapter 4. Basic JPA CRUD Commands

JPA provides an API for implementing persistence to the database through manipulation of `@Entity` instances and calls to the `EntityManager`.

## 4.1. EntityManager persist()

We create a new object in the database by calling `persist()` on the `EntityManager` and passing in an `@Entity` instance that represents something new. This will:

- assign a primary key if configured to do so
- add the instance to the Persistence Context
- make the `@Entity` instance managed from that point forward

The following snippet shows a partial DAO implementation using JPA.

*Example EntityManager persist() Call*

```
@Component
@RequiredArgsConstructor
public class JpaSongDAO {
    private final EntityManager em; ①

    public void create(Song song) {
        em.persist(song); ②
    }
    ...
}
```

① using an injected `EntityManager`; it is important that it be injected

② stage for insertion into database; all instance changes managed from this point forward

A database `INSERT` SQL command will be queued to the database as a result of a successful call and the `@Entity` instance will be in a managed state.

*Resulting SQL from persist Call()*

```
Hibernate: call next value for reposongs_song_sequence
Hibernate: insert into reposongs_song (artist, released, title, id) values (?, ?, ?, ?)
binding parameter (1:VARCHAR) <- [The Orb]
binding parameter (2:DATE) <- [2020-08-01]
binding parameter (3:VARCHAR) <- [Mother Night teal]
binding parameter (4:INTEGER) <- [1]
```

In the managed state, any changes to the `@Entity` will result in either the changes be part of:

- a delayed future `INSERT`



- a future **UPDATE** SQL command if changes occur after the **INSERT**

Updates are issued during the next JPA session "flush". JPA session flushes can be triggered manually or automatically prior to or no later than the next commit.

## 4.2. EntityManager find() By Identity

JPA supplies a means to get the full **@Entity** using its primary key.

*Example EntityManager find() Call*

```
public Song findById(int id) {
    return em.find(Song.class, id);
}
```

If the instance is not yet loaded into the Persistence Context, **SELECT** SQL command(s) will be issued to the database to obtain the persisted state. The following snippet shows the SQL generated by Hibernate to fetch the state from the database to realize the **@Entity** instance within the JVM.

*Resulting SQL from find() Call*

```
Hibernate: select
    s1_0.id,
    s1_0.artist,
    s1_0.released,
    s1_0.title
from reposongs_song s1_0
where s1_0.id=?
```

From that point forward, the state will be returned from the Persistence Context without the need to get the state from the database.



Find by identity is resolved first through the local cache and second through a database query. Once the **@Entity** is managed, all explicit find() by identity calls will be resolved without a database query. That will not be true when using a query (discussed next) with caller-provided criteria.

## 4.3. EntityManager query

JPA provides many types of queries

- JPA Query Language (officially abbreviated "JPQL"; often called "JPAQL") - a very SQL-like String syntax expressed in terms of **@Entity** classes and relationship constructs
- Criteria Language - a type-safe, Java-centric syntax that avoids String parsing and makes dynamic query building more efficient than query string concatenation and parsing
- Native SQL - the same SQL we would have provided to JDBC

The following snippet shows an example of executing a JPQL Query.

#### Example EntityManager Query

```
public boolean existsById(int id) {
    return em.createQuery("select count(s) from Song s where s.id=:id", ①
        Number.class) ②
        .setParameter("id", id) ③
        .getSingleResult() ④
        .longValue()==1L; ⑤
}
```

- ① JPQL String based on `@Entity` constructs
- ② query call syntax allows us to define the expected return type
- ③ query variables can be set by name or position
- ④ one (mandatory) or many results can be returned from a query
- ⑤ entity exists if row count of rows matching PK is 1. Otherwise, should be 0

The following shows how our JPQL snippet mapped to the raw SQL issued to the database. Notice that our `Song @Entity` reference was mapped to the `REPOSONGS_SONG` database table.

#### Resulting SQL from Query Call

```
Hibernate: select
    count(s1_0.id)
  from reposongs_song s1_0
 where s1_0.id=?
```

## 4.4. EntityManager flush()

Not every change to an `@Entity` and call to an `EntityManager` results in an immediate 1:1 call to the database. Some of these calls manipulate an in-memory cache in the JVM and may get issued in a group of other commands at some point in the future. We normally want to allow the `EntityManager` to cache these calls as much as possible. However, there are times (e.g., before making a raw SQL query) where we want to make sure the database has the current state of the cache.

The following snippet shows an example of flushing the contents of the cache after changing the state of a managed `@Entity` instance.

#### Example EntityManager flush() Call

```
Song s = ... //obtain a reference to a managed instance
s.setTitle("...");
em.flush(); //optional!!! will eventually happen at some point
```

Whether it was explicitly issued or triggered internally by the JPA provider, the following snippet shows the resulting `UPDATE` SQL call to change the state of the database to match the Persistence

Context.

*Resulting SQL from flush() Call*

```
Hibernate: update reposongs_song
  set artist=?, released=?, title=? ①
  where id=?
```

① all fields designated as `updatable=true` are included in the `UPDATE`



*flush() does not Commit Changes*

Flushing commands to the database only makes the changes available to the current transaction. The result of the commands will not be available to console queries and future transactions until the changes have been committed.

## 4.5. EntityManager remove()

JPA provides a means to delete an `@Entity` from the database. However, we must have the managed `@Entity` instance loaded in the Persistence Context first to use this capability. The reason for this is that a JPA delete can optionally involve cascading actions to remove other related entities as well.

The following snippet shows how a managed `@Entity` instance can be used to initiate the removal from the database.

*Example EntityManager remove() Call*

```
public void delete(Song song) {
    em.remove(song);
}
```

The following snippet shows how the remove command was mapped to a SQL `DELETE` command.

*Resulting SQL from remove() Call*

```
Hibernate: delete from reposongs_song
  where id=?
```

## 4.6. EntityManager clear() and detach()

There are two commands that will remove entities from the Persistence Context. They have their purpose, but know that they are rarely used and can be dangerous to call.

- `clear()` - will remove all entities
- `detach()` - will remove a specific `@Entity`

I only bring these up because you may come across class examples where I am calling `flush()` and `clear()` in the middle of a demonstration. This is purposely mimicking a fresh Persistence Context

within scope of a single transaction.

#### *clear() and detach() Commands*

```
em.clear();  
em.detach(song);
```

Calling `clear()` or `detach()` will evict all managed entities or targeted managed `@Entity` from the Persistence Context—losing any in-progress and future modifications. In the case of returning redacted `@Entities`—this may be exactly what you want (you don't want the redactions to remove data from the database).



#### *Use clear() and detach() with Caution*

Calling `clear()` or `detach()` will evict all managed entities or targeted managed `@Entity` from the Persistence Context—losing any in-progress and future modifications.

# Chapter 5. Transactions

All commands require some type of transaction when interacting with the database. The transaction can be activated and terminated at varying levels of scope integrating one or more commands into a single transaction.

## 5.1. Transactions Required for Explicit Changes/Actions

The injected `EntityManager` is the target of our application calls, and the transaction gets associated with that object. The following snippet shows the provider throwing a `TransactionRequiredException` when the calling `persist()` on the injected `EntityManager` when no transaction has been activated.

*Example Persist Failure without Transaction*

```
@Autowired
private EntityManager em;
...
@Test
void transaction_missing() {
    //given - an instance
    Song song = mapper.map(dtoFactory.make());

    //when - persist is called without a tx, an exception is thrown
    em.persist(song); ❶
}
```

❶ `TransactionRequiredException` exception thrown

*Exception Thrown when Required Transaction Missing*

```
jakarta.persistence.TransactionRequiredException: No EntityManager with actual
transaction available for current thread - cannot reliably process 'persist' call
```

## 5.2. Activating Transactions

Although you will find transaction methods on the `EntityManager`, these are only meant for individually managed instances created directly from the `EntityManagerFactory`. Transactions for injected an `EntityManager` are managed by the container and triggered by the presence of a `@Transactional` annotation on a called bean method within the call stack.

This next example annotates the calling `@Test` method with the `@Transactional` annotation to cause a transaction to be active for the three (3) contained `EntityManager` calls.

*Example @Transactional Activation*

```
import org.springframework.transaction.annotation.Transactional;
```

```

...
@Test
@Transactional ①
void transaction_present_in_caller() {
    //given - an instance
    Song song = mapper.map(dtoFactory.make());

    //when - persist called within caller transaction, no exception thrown
    em.persist(song); ②
    em.flush(); //force DB interaction ②

    //then
    then(em.find(Song.class, song.getId())).isNotNull(); ②
} ③

```

- ① **@Transactional** triggers an Aspect to activate a transaction for the Persistence Context operating within the current thread
- ② the same transaction is used on all three (3) **EntityManager** calls
- ③ the end of the method will trigger the transaction-initiating Aspect to commit (or rollback) the transaction it activated

Nested calls annotated with **@Transactional**, by default, will continue the current transaction.

## 5.3. Conceptual Transaction Handling

Logically speaking, the transaction handling done on behalf of **@Transactional** is similar to the snippet shown below. However, as complicated as that is — it does not begin to address nested calls. Also note that a thrown **RuntimeException** triggers a rollback and anything else triggers a commit.

### *Conceptual View of Transaction Handling*

```

tx = em.getTransaction();
try {
    tx.begin();
    //call code ②
} catch (RuntimeException ex) {
    tx.setRollbackOnly(); ①
} finally { ②
    if (tx.getRollbackOnly()) {
        tx.rollback();
    } else {
        tx.commit();
    }
}

```

- ① **RuntimeException**, by default, triggers a rollback
- ② Normal returns and checked exceptions, by default, trigger a commit

## 5.4. Activating Transactions in @Components

We can alternatively push the demarcation of the transaction boundary down to the `@Component` methods.

The snippet below shows a DAO `@Component` that designates each of its methods being `@Transactional`. This has the benefit of knowing that each of the calls to `EntityManager` methods will have the required transaction in place. Whether a new one is the right one is a later topic.

*@Transactional Component*

```
@Component
@RequiredArgsConstructor
@Transactional ①
public class JpaSongDAO {
    private final EntityManager em;

    public void create(Song song) {
        em.persist(song);
    }
    public Song findById(int id) {
        return em.find(Song.class, id);
    }
    public void delete(Song song) {
        em.remove(song);
    }
}
```

① each method will be assigned to a transaction

## 5.5. Calling @Transactional @Component Methods

The following example shows the calling code invoking methods of the DAO `@Component` in independent transactions. The code works because there really is no dependency between the `INSERT` and `SELECT` to be part of the same transaction, as long as the `INSERT` commits before the `SELECT` transaction starts.

*Calling @Component @Transactional Methods*

```
@Test
void transaction_present_in_component() {
    //given - an instance
    Song song = mapper.map(dtoFactory.make());

    //when - persist called within component transaction, no exception thrown
    jpaDao.create(song); ①

    //then
    then(jpaDao.findById(song.getId())).isNotNull(); ②
}
```

① **INSERT** is completed in a separate transaction

② **SELECT** completes in follow-on transaction

## 5.6. @Transactional @Component Methods SQL

The following shows the SQL triggered by the snippet above with the different transactions annotated.

*@Transactional Methods Resulting SQL*

```
①
Hibernate: insert into reposongs_song (artist,released,title,id) values (?,?,,?)
②
Hibernate: select
    s1_0.id,
    s1_0.artist,
    s1_0.released,
    s1_0.title
from reposongs_song s1_0
where s1_0.id=?
```

① transaction 1

② transaction 2

## 5.7. Unmanaged @Entity

However, we do not always get that lucky—for individual, sequential transactions to play well together. JPA entities follow the notation of managed and unmanaged/detached state.

- Managed entities are actively being tracked by a Persistence Context
- Unmanaged/Detached entities have either never been or no longer associated with a Persistence Context

The following snippet shows an example of where a follow-on method fails because the **EntityManager** requires that **@Entity** be currently managed. However, the end of the **create()** transaction made it detached.

*Unmanaged @Entity*

```
@Test
void transaction_common_needed() {
    //given a persisted instance
    Song song = mapper.map(dtoFactory.make());
    jpaDao.create(song); //song is detached at this point ①

    //when - removing detached entity we get an exception
    jpaDao.delete(song); ②
```



- ① the first transaction starts and ends at this call
- ② the `EntityManager.remove` operates in a separate transaction with a detached `@Entity` from the previous transaction

The following text shows the error message thrown by the `EntityManager.remove` call when a detached entity is passed in to be deleted.

```
java.lang.IllegalArgumentException: Removing a detached instance
info.ejava.examples.db.repo.jpa.songs.bo.Song#1
```

## 5.8. Shared Transaction

We can get things to work better if we encapsulate methods behind a `@Service` method defining good transaction boundaries. Lacking a more robust application, the snippet below adds the `@Transactional` to the `@Test` method to have it shared by the three (3) DAO `@Component` calls — making the `@Transactional` annotations on the DAO meaningless.

*Shared Transaction*

```
@Test
@Transactional ①
void transaction_common_present() {
    //given a persisted instance
    Song song = mapper.map(dtoFactory.make());
    jpaDao.create(song); //song is detached at this point ②

    //when - removing managed entity, it works
    jpaDao.delete(song); ②

    //then
    then(jpaDao.findById(song.getId())).isNull(); ②
}
```

- ① `@Transactional` at the calling method level is shared across all lower-level calls
- ② Each DAO call is executed in the same transaction and the `@Entity` can still be managed across all calls

## 5.9. @Transactional Attributes

There are several attributes that can be set on the `@Transactional` annotation. A few of the more common properties to set include

- propagation - defaults to REQUIRED, proactively activating a transaction if not already present
  - SUPPORTS - lazily initiates a transaction, but fully supported if already active
  - MANDATORY - error if called without an active transaction
  - REQUIRES\_NEW - proactively creates a new transaction separate from the caller's

transaction

- NOT\_SUPPORTED - nothing within the called method will honor transaction semantics
- NEVER - do not call with an active transaction
- NESTED - may not be supported, but permits nested transactions to complete before returning to calling transaction
- isolation - location to assign JDBC Connection isolation
- readOnly - defaults to false; hints to JPA provider that entities can be immediately detached
- rollback definitions - when to implement non-standard rollback rules

# Chapter 6. Summary

In this module, we learned:

- to configure a JPA project, include project dependencies, and required application properties
- to define a Persistence Context and where to scan for `@Entity` classes
- requirements for an `@Entity` class
- default mapping conventions for `@Entity` mappings
- optional mapping annotations for `@Entity` mappings
- to perform basic CRUD operations with the database