



Pro JPA 2 in Java EE 8

An In-Depth Guide to Java Persistence APIs

—

Third Edition

—

Mike Keith
Merrick Schincariol
Massimo Nardone

Apress®

Pro JPA 2 in Java EE 8

**An In-Depth Guide to Java
Persistence APIs**

Third Edition

Mike Keith

Merrick Schincariol

Massimo Nardone

Apress®

Pro JPA 2 in Java EE 8: An In-Depth Guide to Java Persistence APIs

Mike Keith
Ottawa, Ontario, Canada

Merrick Schincariol
Almonte, Ontario, Canada

Massimo Nardone
Helsinki, Finland

ISBN-13 (pbk): 978-1-4842- 3419-8
<https://doi.org/10.1007/978-1-4842-3420-4>

ISBN-13 (electronic): 978-1-4842- 3420-4

Library of Congress Control Number: 2018932342

Copyright © 2018 by Mike Keith, Merrick Schincariol, Massimo Nardone

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Cover image by Freepik (www.freepik.com)

Managing Director: Welmoed Spahr
Editorial Director: Todd Green
Acquisitions Editor: Steve Anglin
Development Editor: Matthew Moodie
Technical Reviewer: Mario Faliero
Coordinating Editor: Mark Powers
Copy Editor: Kezia Endsley

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit <http://www.apress.com/rights-permissions>.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/9781484234198. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

To my wife Darleen, the perfect mother, and to Cierra, Ariana, Jeremy, and Emma, who brighten my life and make me strive to be a better person.

—Mike

To Anthony, whose boundless creativity continues to inspire me. To Evan, whose boisterous enthusiasm motivates me to take on new challenges. To Kate, who proves that size is no object when you have the right attitude. I love you all.

—Merrick

I would like to dedicate this book to the memory of my beloved late mother Maria Augusta Ciniglio. Thanks mom for all the great things you taught me, for making me a good person, for making me study to become a computing scientist, and for the great memories you left me. You will be loved and missed forever. I love you mom. RIP.

—Massimo

Table of Contents

- About the Authors.....xvii
- About the Technical Reviewerxix
- Acknowledgmentsxxi
- Chapter 1: Introduction..... 1
 - Relational Databases..... 2
 - Object-Relational Mapping..... 3
 - The Impedance Mismatch 4
 - Java Support for Persistence 11
 - Proprietary Solutions 11
 - JDBC 13
 - Enterprise JavaBeans 13
 - Java Data Objects 15
 - Why Another Standard?..... 16
 - The Java Persistence API 17
 - History of the Specification 17
 - Overview 21
 - Summary..... 24
- Chapter 2: Getting Started 25
 - Entity Overview 25
 - Persistability 26
 - Identity..... 26
 - Transactionality 27
 - Granularity 27

TABLE OF CONTENTS

Entity Metadata	28
Annotations.....	28
XML.....	30
Configuration by Exception	30
Creating an Entity	31
Entity Manager	34
Obtaining an Entity Manager	36
Persisting an Entity.....	37
Finding an Entity.....	38
Removing an Entity.....	39
Updating an Entity	40
Transactions	41
Queries	42
Putting It All Together	44
Packaging It Up	47
Persistence Unit.....	47
Persistence Archive	48
Summary	49
Chapter 3: Enterprise Applications.....	51
Application Component Models.....	54
Session Beans	56
Stateless Session Beans	57
Stateful Session Beans.....	61
Singleton Session Beans	65
Servlets	67
Dependency Management and CDI	69
Dependency Lookup	70
Dependency Injection	72
Declaring Dependencies.....	74

CDI and Contextual Injection	78
CDI Beans	78
Injection and Resolution	79
Scopes and Contexts	80
Qualified Injection.....	81
Producer Methods and Fields.....	82
Using Producer Methods with JPA Resources.....	83
Transaction Management.....	85
Transaction Review	85
Enterprise Transactions in Java	86
Putting It All Together	95
Defining the Component.....	96
Defining the User Interface.....	97
Packaging It Up.....	98
Summary	99
Chapter 4: Object-Relational Mapping.....	101
Persistence Annotations.....	102
Accessing Entity State.....	103
Field Access.....	104
Property Access.....	105
Mixed Access.....	106
Mapping to a Table	108
Mapping Simple Types	110
Column Mappings.....	111
Lazy Fetching.....	113
Large Objects.....	114
Enumerated Types	115
Temporal Types	118
Transient State.....	119

TABLE OF CONTENTS

- Mapping the Primary Key 120
 - Overriding the Primary Key Column 120
 - Primary Key Types 121
 - Identifier Generation 121
- Relationships..... 129
 - Relationship Concepts 129
 - Mappings Overview 132
 - Single-Valued Associations 133
 - Collection-Valued Associations..... 140
 - Lazy Relationships..... 148
- Embedded Objects 149
- Summary 154
- Chapter 5: Collection Mapping 157**
 - Relationships and Element Collections 157
 - Using Different Collection Types..... 161
 - Sets or Collections..... 162
 - Lists 162
 - Maps..... 167
 - Duplicates..... 185
 - Null Values 187
 - Best Practices 188
 - Summary 189
- Chapter 6: Entity Manager 191**
 - Persistence Contexts..... 191
 - Entity Managers 192
 - Container-Managed Entity Managers 192
 - Application-Managed Entity Managers..... 198
 - Transaction Management..... 201
 - JTA Transaction Management..... 202
 - Resource-Local Transactions..... 218
 - Transaction Rollback and Entity State 221

Choosing an Entity Manager	224
Entity Manager Operations.....	225
Persisting an Entity.....	225
Finding an Entity.....	227
Removing an Entity.....	228
Cascading Operations.....	229
Clearing the Persistence Context	234
Synchronization with the Database.....	234
Detachment and Merging.....	238
Detachment	238
Merging Detached Entities	241
Working with Detached Entities.....	246
Summary.....	267
Chapter 7: Using Queries	269
Java Persistence Query Language	270
Getting Started	270
Filtering Results.....	271
Projecting Results.....	272
Joins Between Entities	272
Aggregate Queries	273
Query Parameters.....	273
Defining Queries.....	274
Dynamic Query Definition	275
Named Query Definition.....	278
Dynamic Named Queries	280
Parameter Types.....	282
Executing Queries	285
Working with Query Results	287
Stream Query Results.....	288
Query Paging	293

TABLE OF CONTENTS

Queries and Uncommitted Changes	296
Query Timeouts.....	299
Bulk Update and Delete.....	300
Using Bulk Update and Delete	301
Bulk Delete and Relationships.....	304
Query Hints.....	305
Query Best Practices	307
Named Queries	307
Report Queries.....	308
Vendor Hints	308
Stateless Beans	309
Bulk Update and Delete	309
Provider Differences	310
Summary	310
Chapter 8: Query Language	313
Introducing JP QL	313
Terminology	314
Example Data Model.....	315
Example Application	316
Select Queries	319
SELECT Clause.....	321
FROM Clause	325
WHERE Clause	336
Inheritance and Polymorphism.....	344
Scalar Expressions	347
ORDER BY Clause	353
Aggregate Queries.....	354
Aggregate Functions.....	356
GROUP BY Clause	357
HAVING Clause.....	358
Update Queries.....	359

Delete Queries	360
Summary	361
Chapter 9: Criteria API	363
Overview	363
The Criteria API	364
Parameterized Types	365
Dynamic Queries	366
Building Criteria API Queries	370
Creating a Query Definition	370
Basic Structure	372
Criteria Objects and Mutability	373
Query Roots and Path Expressions	374
The SELECT Clause	377
The FROM Clause	382
The WHERE Clause	384
Building Expressions	385
The ORDER BY Clause	401
The GROUP BY and HAVING Clauses	402
Bulk Update and Delete	403
Strongly Typed Query Definitions	405
The Metamodel API	405
Strongly Typed API Overview	407
The Canonical Metamodel	409
Choosing the Right Type of Query	412
Summary	413
Chapter 10: Advanced Object-Relational Mapping	415
Table and Column Names	416
Converting Entity State	418
Creating a Converter	418
Declarative Attribute Conversion	420

TABLE OF CONTENTS

Automatic Conversion.....	423
Converters and Queries	424
Complex Embedded Objects.....	425
Advanced Embedded Mappings	425
Overriding Embedded Relationships.....	427
Compound Primary Keys	429
ID Class.....	430
Embedded ID Class.....	432
Derived Identifiers	434
Basic Rules for Derived Identifiers	435
Shared Primary Key.....	436
Multiple Mapped Attributes	439
Using EmbeddedId.....	440
Advanced Mapping Elements.....	443
Read-Only Mappings	444
Optionality	445
Advanced Relationships	446
Using Join Tables.....	446
Avoiding Join Tables	447
Compound Join Columns.....	449
Orphan Removal	451
Mapping Relationship State.....	453
Multiple Tables	456
Inheritance	461
Class Hierarchies.....	461
Inheritance Models	466
Mixed Inheritance.....	477
Summary.....	480

Chapter 11: Advanced Queries	483
SQL Queries.....	483
Native Queries vs. JDBC	484
Defining and Executing SQL Queries	487
SQL Result Set Mapping	491
Parameter Binding	500
Stored Procedures	500
Entity Graphs	505
Entity Graph Annotations	507
Entity Graph API	516
Managing Entity Graphs	519
Using Entity Graphs	522
Summary	525
Chapter 12: Other Advanced Topics.....	527
Lifecycle Callbacks.....	527
Lifecycle Events.....	528
Callback Methods	529
Entity Listeners	531
Inheritance and Lifecycle Events.....	536
Validation.....	542
Using Constraints	543
Invoking Validation.....	545
Validation Groups.....	546
Creating New Constraints	548
Validation in JPA	551
Enabling Validation	552
Setting Lifecycle Validation Groups	553

TABLE OF CONTENTS

Concurrency	555
Entity Operations	555
Entity Access	555
Refreshing Entity State.....	555
Locking.....	559
Optimistic Locking	560
Pessimistic Locking	574
Caching	580
Sorting Through the Layers	580
Shared Cache	582
Utility Classes	589
PersistenceUtil.....	589
PersistenceUnitUtil	590
Summary	591
Chapter 13: XML Mapping Files.....	593
The Metadata Puzzle	595
The Mapping File.....	596
Disabling Annotations	598
Persistence Unit Defaults	601
Mapping File Defaults.....	606
Queries and Generators	609
Managed Classes and Mappings.....	617
Converters	652
Summary	654
Chapter 14: Packaging and Deployment.....	655
Configuring Persistence Units	656
Persistence Unit Name	656
Transaction Type	657
Persistence Provider.....	658
Data Source	659

Mapping Files	662
Managed Classes	663
Shared Cache Mode.....	667
Validation Mode	668
Adding Properties	668
Building and Deploying	669
Deployment Classpath.....	669
Packaging Options.....	670
Persistence Unit Scope.....	676
Outside the Server.....	677
Configuring the Persistence Unit.....	678
Specifying Properties at Runtime	680
System Classpath	681
Schema Generation	682
The Generation Process.....	683
Deployment Properties	684
Runtime Properties.....	689
Mapping Annotations Used by Schema Generation	689
Unique Constraints	690
Null Constraints	691
Indexes	692
Foreign Key Constraints.....	692
String-Based Columns.....	694
Floating Point Columns.....	695
Defining the Column	696
Summary	697
Chapter 15: Testing.....	699
Testing Enterprise Applications	699
Terminology	700
Testing Outside the Server	702
JUnit	703

TABLE OF CONTENTS

Unit Testing..... 704

 Testing Entities 704

 Testing Entities in Components 706

 The Entity Manager in Unit Tests 709

Integration Testing..... 713

 Using the Entity Manager 714

 Components and Persistence 722

 Test Frameworks 736

Best Practices 738

Summary 739

Index..... 741

About the Authors



Mike Keith was the co-specification lead for JPA 1.0 and a member of the JPA 2.0 and JPA 2.1 expert groups. He sits on a number of other Java Community Process expert groups and the Enterprise Expert Group (EEG) in the OSGi Alliance. He holds a Master's degree in Computer Science from Carleton University, and has over 20 years experience in persistence and distributed systems research and practice. He has written papers and articles on JPA and spoken at numerous conferences around the world. He is employed as an architect at Oracle in Ottawa, Canada, and is married with four kids and two dogs.



Merrick Schincariol is a consulting engineer at Oracle, specializing in middleware technologies. He has a Bachelor of Science degree in Computer Science from Lakehead University, and has more than a decade of experience in enterprise software development. He spent some time consulting in the pre-Java enterprise and business intelligence fields before moving on to write Java and J2EE applications. His experience with large-scale systems and data warehouse design gave him a mature and practiced perspective on enterprise software, which later propelled him into doing Java EE container implementation work.

ABOUT THE AUTHORS



Massimo Nardone has more than 24 years of experience in Security, Web/mobile development, cloud, and IT architecture. His true IT passions are security and Android.

He has been programming and teaching others how to program with Android, Perl, PHP, Java, VB, Python, C/C++, and MySQL for more than 20 years.

He holds a Master of Science degree in Computing Science from the University of Salerno, Italy. He has worked as a Project Manager, Software Engineer, Research Engineer, Chief Security Architect, Information Security Manager,

PCI/SCADA Auditor, and Senior Lead IT Security/Cloud/SCADA Architect for many years.

His technical skills include security, Android, cloud, Java, MySQL, Drupal, Cobol, Perl, Web and mobile development, MongoDB, D3, Joomla, Couchbase, C/C++, WebGL, Python, Pro Rails, Django CMS, Jekyll, Scratch, and more.

He worked as a visiting lecturer and supervisor at the Networking Laboratory of the Helsinki University of Technology (Aalto University). He also holds four international patents (in the PKI, SIP, SAML, and Proxy areas).

He currently works as Chief Information Security Office (CISO) for Cargotec Oyj and is a member of the ISACA Finland Chapter Board.

Massimo has reviewed more than 45 IT books for different publishers and is the coauthor of *Pro Android Games* (Apress, 2015).

About the Technical Reviewer



Mario Faliero is a telecommunications engineer and entrepreneur. He has more than ten years of experience with radio frequency hardware engineering. Mario has extensive experience in numerical coding, using scripting languages (MATLAB and Python) and compiled languages (C/C++ and Java). He has been responsible for the development of electromagnetic assessment tools for space and commercial applications. Mario received his Master's degree from the University of Siena.

Acknowledgments

Many thanks go to my wonderful family—my wife Pia, and my children Luna, Leo, and Neve—for supporting me when working on this book. You are the most beautiful aspect of my life.

I want to thank my beloved late mother Maria Augusta Ciniglio who always supported and loved me so much. I will love and miss you forever, my dearest mom.

I also need to thank my beloved father Giuseppe and my brothers Mario and Roberto for your endless love and for being the best dad and brothers in the world.

This book is also dedicated to Doctor Antonio Catapano, for being such a great person with a big heart and taking care of me and my mother. To my sister in law Susanna Cennamo, to my dear cousins Rosaria Scudieri, Pina and Elisa Franzese, and Francesco Ciniglio, for loving and supporting me and my mother like no other. To Pertti and Marianna Kantola, for teaching me how to be a good programmer, taking care of me, and treating me like their son. To Antti, Piia, and Daniela Jalonen for being great and supportive friends, as well as to Anton Jalonen, who will become a great software engineer. Anton, may this book be an inspiration to your great IT future.

I also want to thank Steve Anglin and Matthew Moodie for giving me the opportunity to write this book. A special thanks goes, as usual, to Mark Powers for doing such a great job and supporting me during the editorial process.

Finally I want to thank Mario Faliero, a good friend and the technical reviewer of this book, for helping me make a better book.

CHAPTER 1

Introduction

Enterprise applications are defined by their need to collect, process, transform, and report on vast amounts of information. And, of course, that information has to be kept somewhere. Storing and retrieving data is a multibillion dollar business, evidenced in part by the growth of the database market as well as the emergence of cloud-based storage services. Despite all the available technologies for data management, application designers still spend much of their time trying to efficiently move their data to and from storage.

Despite the success the Java platform has had in working with database systems, for a long time it suffered from the same problem that has plagued other object-oriented programming languages. Moving data back and forth between a database system and the object model of a Java application was a lot harder than it needed to be. Java developers either wrote lots of code to convert row and column data into objects, or found themselves tied to proprietary frameworks that tried to hide the database from them. Fortunately, a standard solution, the Java Persistence API (JPA), was introduced into the platform to bridge the gap between object-oriented domain models and relational database systems.

This book introduces version 2.2 of the Java Persistence API as part of the Java EE 8 and explores everything that it has to offer developers.

Maintenance release of JPA 2.2 started during 2017 under JSR 338 and was finally approved on June 19, 2017.

Here is the official Java Persistence 2.2 Maintenance release statement:

“The Java Persistence 2.2 specification enhances the Java Persistence API with support for repeating annotations; injection into attribute converters; support for mapping of the `java.time.LocalDate`, `java.time.LocalDateTime`, `java.time.OffsetTime`, and `java.time.OffsetDateTime` types; and methods to retrieve the results of `Query` and `TypedQuery` as streams.”

¹**Electronic supplementary material** The online version of this chapter (https://doi.org/10.1007/978-1-4842-3420-4_1) contains supplementary material, which is available to authorized users.

One of its strengths is that it can be slotted into whichever layer, tier, or framework an application needs it to be in. Whether you are building client-server applications to collect form data in a Swing application or building a website using the latest application framework, JPA can help you provide persistence more effectively.

To set the stage for JPA, this chapter first takes a step back to show where we’ve been and what problems we are trying to solve. From there, we will look at the history of the specification and give you a high-level view of what it has to offer.

Relational Databases

Many ways of persisting data have come and gone over the years, and no concept has more staying power than the relational database. Even in the age of the cloud, when “Big Data” and “NoSQL” regularly steal the headlines, relational database services are in consistent demand to enable today’s enterprise applications running in the cloud. While key-value and document-oriented NoSQL stores have their place, relational stores remain the most popular general-purpose databases in existence, and they are where the vast majority of the world’s corporate data is stored. They are the starting point for every enterprise application and often have a lifespan that continues long after the application has faded away.

Understanding relational data is key to successful enterprise development. Developing applications to work well with database systems is a commonly acknowledged hurdle of software development. A good deal of Java’s success can be attributed to its widespread adoption for building enterprise database systems. From consumer websites to automated gateways, Java applications are at the heart of enterprise application development. Figure 1-1 shows an example of a relational database of user to car.

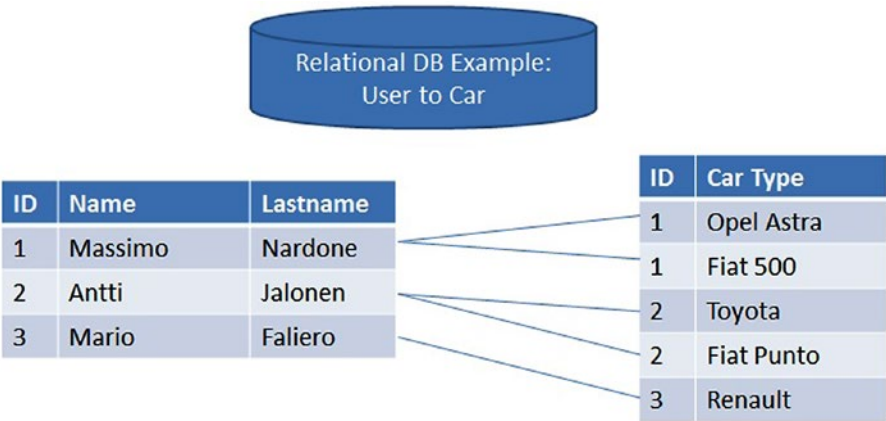


Figure 1-1. User to car relational database

Object-Relational Mapping

“The domain model has a class. The database has a table. They look pretty similar. It should be simple to convert one to the other automatically.” This is a thought we’ve probably all had at one point or another while writing yet another Data Access Object (DAO) to convert Java Database Connectivity (JDBC) result sets into something object-oriented. The domain model looks similar enough to the relational model of the database that it seems to cry out for a way to make the two models talk to each other.

The technique of bridging the gap between the object model and the relational model is known as object-relational mapping, often referred to as O-R mapping or simply ORM. The term comes from the idea that we are in some way mapping the concepts from one model onto another, with the goal of introducing a mediator to manage the automatic transformation of one to the other.

Before going into the specifics of object-relational mapping, let’s define a brief manifesto of what the ideal solution should be.

- *Objects, not tables:* Applications should be written in terms of the domain model, not bound to the relational model. It must be possible to operate on and query against the domain model without having to express it in the relational language of tables, columns, and foreign keys.
- *Convenience, not ignorance:* Mapping tools should be used only by someone familiar with relational technology. O-R mapping is not meant to save developers from understanding mapping problems or to hide them altogether. It is meant for those who have an understanding of the issues and know what they need, but who don’t want to have to write thousands of lines of code to deal with a problem that has already been solved.
- *Unobtrusive, not transparent:* It is unreasonable to expect that persistence be transparent because an application always needs to have control of the objects that it is persisting and be aware of the entity lifecycle. The persistence solution should not intrude on the domain model, however, and domain classes must not be required to extend classes or implement interfaces in order to be persistable.

- *Legacy data, new objects*: It is far more likely that an application will target an existing relational database schema than create a new one. Support for legacy schemas is one of the most relevant use cases that will arise, and it is quite possible that such databases will outlive every one of us.
- *Enough, but not too much*: Enterprise developers have problems to solve, and they need features sufficient to solve those problems. What they don't like is being forced to eat a heavyweight persistence model that introduces large overhead because it is solving problems that many do not even agree *are* problems.
- *Local, but mobile*: A persistent representation of data does not need to be modeled as a full-fledged remote object. Distribution is something that exists as part of the application, not part of the persistence layer. The entities that contain the persistent state, however, must be able to travel to whichever layer needs them so that if an application is distributed, then the entities will support and not inhibit a particular architecture.
- *Standard API, with pluggable implementations*: Large companies with sizable applications don't want to risk being coupled to product-specific libraries and interfaces. By depending only on defined standard interfaces, the application is decoupled from proprietary APIs and can switch implementations if another becomes more suitable.

This would appear to be a somewhat demanding set of requirements, but it is one born of both practical experience and necessity. Enterprise applications have very specific persistence needs, and this shopping list of items is a fairly specific representation of the experience of the enterprise community.

The Impedance Mismatch

Advocates of object-relational mapping often describe the difference between the object model and the relational model as the impedance mismatch between the two. This is an apt description because the challenge of mapping one to the other lies not in the similarities between the two, but in the concepts in each for which there is no logical equivalent in the other.

In the following sections, we present some basic object-oriented domain models and a variety of relational models to persist the same set of data. As you will see, the challenge in object-relational mapping is not so much the complexity of a single mapping but that there are so many possible mappings. The goal is not to explain how to get from one point to the other, but to understand the roads that may have to be taken to arrive at an intended destination.

Class Representation

Let’s begin this discussion with a simple class. Figure 1-2 shows an Employee class with four attributes: employee ID, employee name, start date, and current salary.

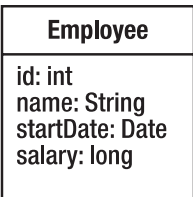


Figure 1-2. The Employee class

Now consider the relational model shown in Figure 1-3. The ideal representation of this class in the database corresponds to scenario (A). Each field in the class maps directly to a column in the table. The employee ID becomes the primary key. With the exception of some slight naming differences, this is a straightforward mapping.

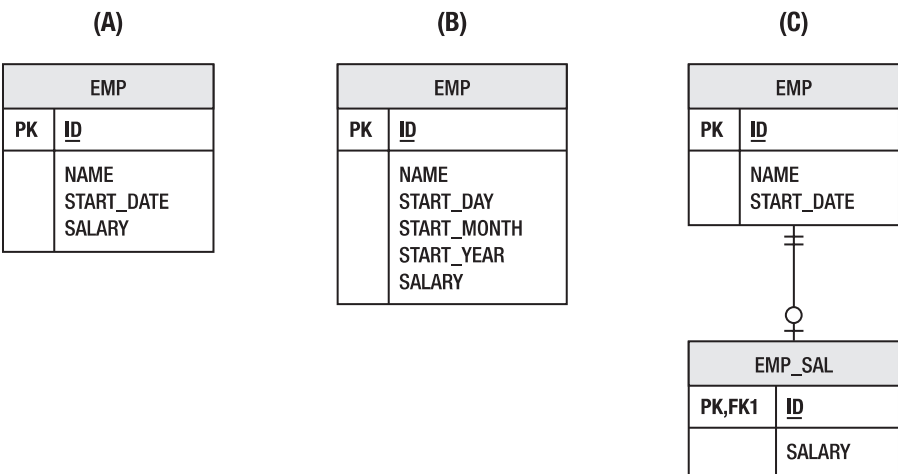


Figure 1-3. Three scenarios for storing employee data

In scenario (B), we see that the start date of the employee is actually stored as three separate columns, one each for the day, month, and year. Recall that the class used a `Date` object to represent this value. Because database schemas are much harder to change, should the class be forced to adopt the same storage strategy in order to remain consistent with the relational model? Also consider the inverse of the problem, in which the class had used three fields, and the table used a single date column. Even a single field becomes complex to map when the database and object model differ in representation.

Salary information is considered commercially sensitive, so it may be unwise to place the salary value directly in the `EMP` table, which may be used for a number of purposes. In scenario (C), the `EMP` table has been split so that the salary information is stored in a separate `EMP_SAL` table. This allows the database administrator to restrict `SELECT` access on salary information to those users who genuinely require it. With such a mapping, even a single store operation for the `Employee` class now requires inserts or updates to two different tables.

Clearly, even storing the data from a single class in a database can be a challenging exercise. We concern ourselves with these scenarios because real database schemas in production systems were never designed with object models in mind. The rule of thumb in enterprise applications is that the needs of the database trump the wants of the application. In fact, there are usually many applications, some object-oriented and some based on Structured Query Language (SQL), that retrieve from and store data into a single database. The dependency of multiple applications on the same database means that changing the database would affect every one of the applications, clearly an undesirable and potentially expensive option. It's up to the object model to adapt and find ways to work with the database schema without letting the physical design overpower the logical application model.

Relationships

Objects rarely exist in isolation. Just like relationships in a database, domain classes depend on and associate themselves with other domain classes. Consider the `Employee` class introduced in Figure 1-2. There are many domain concepts that could be associated with an employee, but for now let's introduce the `Address` domain class, for which an `Employee` may have at most one instance. We say in this case that `Employee` has a one-to-one relationship with `Address`, represented in the Unified Modeling Language (UML) model by the `0..1` notation. Figure 1-4 demonstrates this relationship.

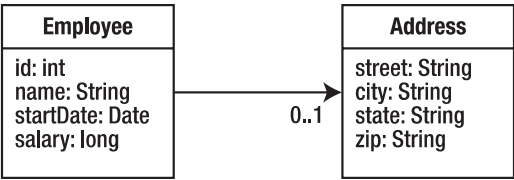


Figure 1-4. *The Employee and Address relationship*

We discussed different scenarios for representing the Employee state in the previous section, and likewise there are several approaches to representing a relationship in a database schema. Figure 1-5 demonstrates three different scenarios for a one-to-one relationship between an employee and an address.

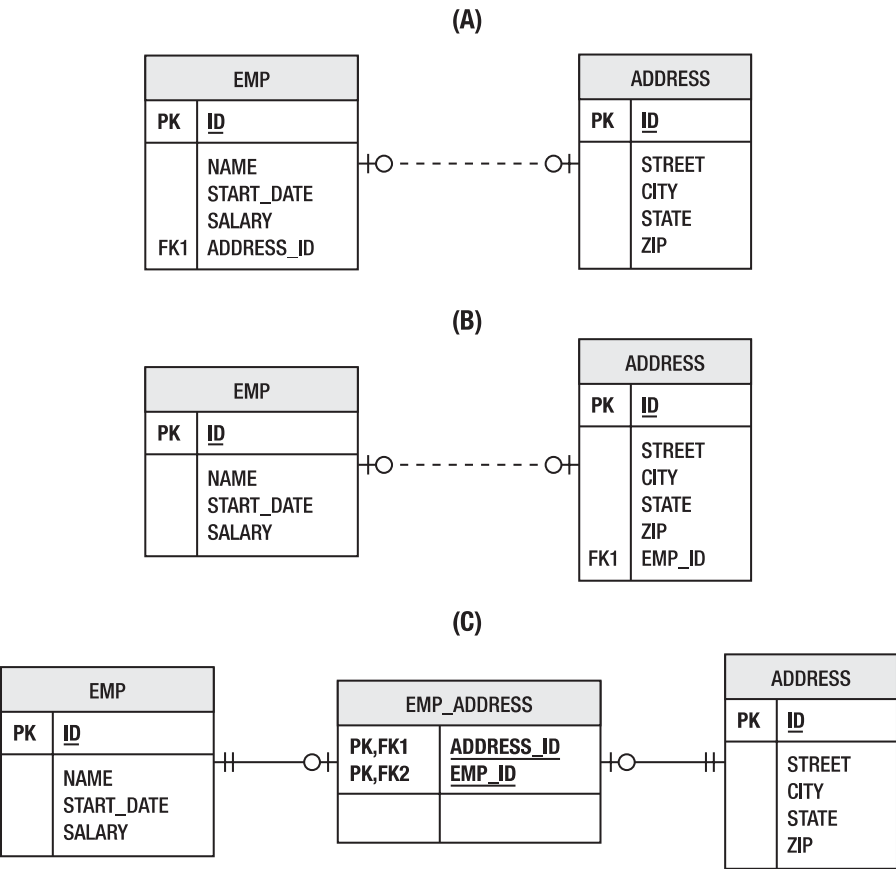


Figure 1-5. *Three scenarios for relating employee and address data*

The building block for relationships in the database is the foreign key. Each scenario involves foreign key relationships between the various tables, but in order for there to be a foreign key relationship, the target table must have a primary key. And so before we even get to associate employees and addresses with each other we have a problem. The domain class `Address` does not have an identifier, yet the table that it would be stored in must have one if it is to be part of relationships. We could construct a primary key out of all of the columns in the `ADDRESS` table, but this is considered bad practice. Therefore, the `ID` column is introduced, and the object relational mapping will have to adapt in some way.

Scenario (A) of Figure 1-5 shows the ideal mapping of this relationship. The `EMP` table has a foreign key to the `ADDRESS` table stored in the `ADDRESS_ID` column. If the `Employee` class holds onto an instance of the `Address` class, the primary key value for the address can be set during store operations when an `EMPLOYEE` row gets written.

And yet consider scenario (B), which is only slightly different yet suddenly much more complex. In the domain model, an `Address` instance did not hold onto the `Employee` instance that owned it, and yet the employee primary key must be stored in the `ADDRESS` table. The object-relational mapping must either account for this mismatch between domain class and table or a reference back to the employee will have to be added for every address.

To make matters worse, scenario (C) introduces a join table to relate the `EMP` and `ADDRESS` tables. Instead of storing the foreign keys directly in one of the domain tables, the join table holds onto the pair of keys. Every database operation involving the two tables must now traverse the join table and keep it consistent. We could introduce an `EmployeeAddress` association class into the domain model to compensate, but that defeats the logical representation we are trying to achieve.

Relationships present a challenge in any object-relational mapping solution. This introduction covered only one-to-one relationships, and yet we have been faced with the need for primary keys not in the object model and the possibility of having to introduce extra relationships into the model or even associate classes to compensate for the database schema.

Inheritance

A defining element of an object-oriented domain model is the opportunity to introduce generalized relationships between like classes. Inheritance is the natural way to express these relationships and allows for polymorphism in the application. Let's revisit the Employee class shown in Figure 1-2 and imagine a company that needs to distinguish between full-time and part-time employees. Part-time employees work for an hourly rate, while full-time employees are assigned a salary. This is a good opportunity for inheritance, moving wage information to the PartTimeEmployee and FullTimeEmployee subclasses. Figure 1-6 shows this arrangement.

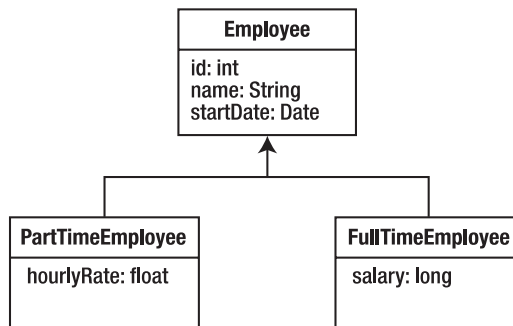


Figure 1-6. *Inheritance relationships between full-time and part-time employees*

Inheritance presents a genuine problem for object-relational mapping. We are no longer dealing with a situation in which there is a natural mapping from a class to a table. Consider the relational models shown in Figure 1-7. Once again, three different strategies for persisting the same set of data are demonstrated.

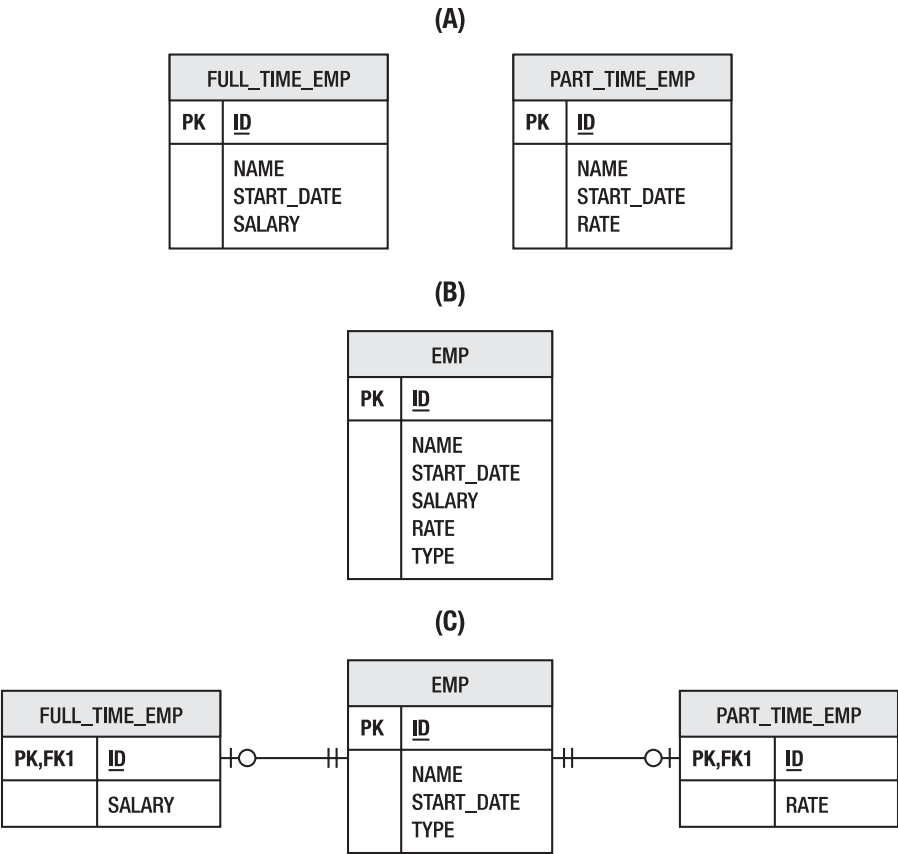


Figure 1-7. *Inheritance strategies in a relational model*

Arguably the easiest solution for someone mapping an inheritance structure to a database would be to put all of the data necessary for each class (including parent classes) into separate tables. This strategy is demonstrated by scenario (A) in Figure 1-7. Note that there is no relationship between the tables (i.e., each table is independent of the others). This means that queries against these tables are now much more complicated if the user needs to operate on both full-time and part-time employees in a single step.

An efficient but denormalized alternative is to place all the data required for every class in the model in a single table. That makes it very easy to query, but note the structure of the table shown in scenario (B) of Figure 1-7. There is a new column, TYPE, which does not exist in any part of the domain model. The TYPE column indicates whether the employee is part-time or full-time. This information must now be interpreted by an object-relational mapping solution to know what kind of domain class to instantiate for any given row in the table.