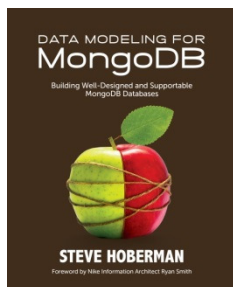# Data Modeling for MongoDB
Building Well-Designed and Supportable MongoDB Databases

Congratulations! You completed the MongoDB application within the given tight timeframe and there is a party to celebrate your application's release into production. Although people are congratulating you at the celebration, you are feeling some uneasiness inside. To complete the project on time required making a lot of assumptions about the data, such as what terms meant and how calculations are derived. In addition, the poor documentation about the application will be of limited use to the support team, and not investigating all of the inherent rules in the data may eventually lead to poorly-performing structures in the not-so-distant future.

Now, what if you had a time machine and could go back and read this book. You would learn that even NoSQL databases like MongoDB require some level of data modeling. **Data modeling is the process of learning about the data, and regardless of technology, this process must be performed for a successful application.** You would learn the value of conceptual, logical, and physical data modeling and how each stage increases our knowledge of the data and reduces assumptions and poor design decisions.

Read this book to learn how to do data modeling for MongoDB applications, and accomplish these five objectives:
1. Understand how data modeling contributes to the process of learning about the data, and is, therefore, a required technique, even when the resulting database is not relational.
2. Know how NoSQL databases differ from traditional relational databases, and where MongoDB fits.
3. Explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts, and learn the basics of adding, querying, updating, and deleting data in MongoDB.
4. Practice a streamlined, template-driven approach to performing conceptual, logical, and physical data modeling. Recognize that *data modeling* does not always have to lead to traditional *data models*!
5. Distinguish top-down from bottom-up development approaches and complete a top-down case study which ties all of the modeling techniques together.

This book is written for anyone who is working with, or will be working with MongoDB, including business analysts, data modelers, database administrators, developers, project managers, and data scientists. There are three sections:

In **Section I**, *Getting Started*, we will reveal the power of data modeling and the tight connections to data models that exist when designing any type of database (Chapter 1), compare NoSQL with traditional relational databases and where MongoDB fits (Chapter 2), explore each MongoDB object and comprehend how each compares to their data modeling and traditional relational database counterparts (Chapter 3), and explain the basics of adding, querying, updating, and deleting data in MongoDB (Chapter 4).

In **Section II**, *Levels of Granularity*, we cover Conceptual Data Modeling (Chapter 5), Logical Data Modeling (Chapter 6), and Physical Data Modeling (Chapter 7). Notice the "ing" at the end of each of these chapters. We focus on the process of building each of these models, which is where we gain essential business knowledge.

In **Section III**, *Case Study*, we will explain both top down and bottom up development approaches and go through a top down case study where we start with business requirements and end with the MongoDB database. This case study will tie together all of the techniques in the previous seven chapters.

Nike Senior Data Architect Ryan Smith wrote the foreword. Key points are included at the end of each chapter as a way to reinforce concepts. In addition, this book is loaded with hands-on exercises, along with their answers provided in Appendix A. Appendix B contains all of the book's references and Appendix C contains a glossary of the terms used throughout the text.