# MongoDB for Jobseekers

*Reach new heights in your career with MongoDB*

**Justin Jenkins**

www.bpbonline.com

First published: 2023

Published by BPB Online
WeWork
119 Marylebone Road
London NW1 5PU

**UK | UAE | INDIA | SINGAPORE**

ISBN 978-93-55518-255

www.bpbonline.com

**Dedicated to**

*My wife* ***Jessica***

*And my sons* ***Caderyn*** *&* ***Evan***

# About the Author

**Justin Jenkins** has nearly two decades of experience in the tech industry, having held various roles. Early on, he found a passion for databases. An early adopter of MongoDB after years of working with SQL, he has made contributions to startups, well known tech companies, and non-profit organizations. Recognized by MongoDB as an Enthusiast, he shares his expertise through resources like this book, online courses which you can find on LinkedIn Learning, and speaking events. Having lived on both coasts of the US, he currently resides in Colorado Springs, Colorado with his wife and two incredibly fun, and incredibly exhausting boys. When not coding, he enjoys cooking, supporting soccer teams like the Seattle Sounders and Manchester United, and exploring the outdoors. He constantly mediates on the words of Mark 9:23-24 and hopes to for the rest of his life.

# About the Reviewers

- **Darshan Jayarama** is a seasoned professional in the field of database administration. With a Bachelor of Engineering degree and over a decade of experience, he has excelled as a MongoDB DBA. Currently serving as a Senior Technical Services Engineer at MongoDB for the past four years, he brings extensive expertise in MongoDB, MySQL, and Sybase. His in-depth knowledge, problem-solving abilities, and commitment to delivering exceptional technical solutions have made him a trusted resource in the industry.

- **Vinicius Grippa** is a Percona Senior Database Engineer, Oracle Ace, and author of the book Learning MySQL. Vinicius has a Bachelor's degree in Computer Science and has worked with databases for over 15 years. He has experience designing databases for mission-critical applications and has become a specialist in MySQL and MongoDB ecosystems.

# Acknowledgement

I am deeply grateful to my family for their unwavering patience and support throughout the process of writing this book.

To my parents, David and Giselle, I extend a special thank you for tolerating that chaotic pile of computer parts in your basement during my formative years. Your belief in my abilities and continuous encouragement have shaped my career in more ways than I can express.

I would like to extend my gratitude to my employers over the years for providing me with opportunities to explore and incorporate MongoDB in various capacities. Their trust in my skills has allowed me to grow both personally and professionally.

A heartfelt appreciation goes out to the team at BPB Publishing for their support and for providing me with the invaluable opportunity to write this book.

It almost goes without saying, but I would also like to aknowlege MongoDB for developing such an exceptional product and I am grateful for the opportunities it has afforded me to share my experiences.

Lastly, I would like to acknowledge the teachers and specialists who supported me during my early years of education, when I faced challenges with reading and writing. Your dedication and belief in my potential have inspired me to overcome obstacles and your impact on my life will always be cherished. I hope that my journey serves as a testament to how much giving a "little extra help" can mean.

# Preface

Within these pages, we invite you on an immersive journey to discover the remarkable capabilities of MongoDB and gain the expertise to harness its full potential. Whether you are an experienced professional seeking to deepen your knowledge or an enthusiastic beginner eager to embrace this cutting-edge technology, this comprehensive guide will be your trusted companion.

As we explore MongoDB, you will uncover its unique qualities that sets it apart from traditional databases. MongoDB's Document Model revolutionizes data organization and retrieval, propelling it to the forefront of the industry.

We will guide you through practical skills, including installation on various operating systems and essential tools such as MongoDB Shell and Compass. Engaging examples will demonstrate how MongoDB's document-based approach efficiently stores and organizes complex data structures.

You will learn to master fundamental operations such as creating, updating, and deleting documents within MongoDB. We will equip you with the skills to work with complex data types as well, using MongoDB's powerful array and embedded document features.

We will delve into crucial concepts such as indexing and collection management, as well as how to enhance query performance and effectively handle data scaling challenges. From Data migration, security, backup strategies, and encryption, we will provide comprehensive guidance to safeguard your data.

Programming with MongoDB is easy and powerful. To show this, we will will explore code examples in Python, Node.JS, and PHP, enabling you to harness MongoDB's capabilities within your applications. We will also discover MongoDB's cloud services, provided by MongoDB Atlas, which offers cloud hosting, database tools, and effortless application development and maintenance.

Finally, we will provide valuable interview preparation resources, bolstering your MongoDB expertise and increasing your chances of success in the job market.

Prepare for an adventure as we unravel the core concepts and features of MongoDB. Whether you aspire to become a MongoDB professional, strengthen your database management skills, or satiate your curiosity, this book is your gateway to unlocking the true potential of MongoDB and your future career opertunties.

*Here is a brief overview of the chapters:*

**Chapter 1: Why MongoDB? –** This chapter will start off with a brief history of databases, and how MongoDB and its Document Modal are different, as well as how these differences make it so powerful and popular.

**Chapter 2: MongoDB Jobs and Roles –** This chapter is dedicated to providing some context and explaination of different job roles available to those with MongoDB experience as well as what sections of this book you might want to pay particular attention to, based on each particular job role.

**Chapter 3: Getting Started –** In this chapter, we will cover installing MongoDB Server on various different operating systems, as well as Docker and MongoDB's Cloud service Atlas. Additionally, we will walk through installing the MongoDB Shell, Tools and the official GUI MongoDB Compass.

**Chapter 4: A Better Way to Store Data – Documents –** In this chapter, we will cover some key details about Documents, rules you should know, and most importantly, some examples of real-life data put into Documents that will help you understand how you might put your own data into Documents.

**Chapter 5: Let's Do It - Create, Update and Delete Documents –** By the end of this chapter, you will have a basic idea of how to create, update and delete documents within MongoDB. We will also discuss how to do each of these actions to multiple documents at one time.

**Chapter 6: Getting What You Want – Querying –** By the end of this chapter, you should be comfortable with querying MongoDB using various query operators, have an understanding of the core concepts of querying with MongoDB, as well as tools that you can use to deal with unique challenges such as case sensitively in MongoDB.

**Chapter 7: Complex Data, Made Simple –** MongoDB's document model allows us to store much more complex data than legacy databases, by using arrays and embedded documents. These data types give us a lot of flexibility, but with flexibility can come complexity. Fortunately, the MongoDB Query API provides robust tools for dealing with these complex types. By the end of this chapter, you should have a solid understanding of how to perform typical queries to find and modify arrays as well as embedded objects. Additionally, you will have a high-level understanding of the many MongoDB operators available for these data types, and how to use them.

**Chapter 8: The MongoDB Aggregation Framework –** For more complex cases, MongoDB offers what is called the Aggregation Framework, which allows a

structured way to formulate a series of steps, called a "pipeline", to get back just the data you need. In this chapter, we will discuss this framework and dive into some examples of its use. Covering the Aggregation Framework in-depth is beyond the scope of this book. However, by the end of this chapter, you should have a solid idea of how the framework works and how you can use it to fit your needs.

**Chapter 9: Planning for Performance - Collections and Indexes –** In this chapter, we will explain what an index is, how MongoDB uses indexes, different index types, different collection types, and how to create, configure and delete indexes and collections. By the end of this chapter, you should have a solid foundation of indexing and collection options in MongoDB and lots of areas you can look further into, if you want to learn more.

**Chapter 10: Getting In and Getting Out - Data Migration –** By the end of this chapter, you should feel comfortable with importing and exporting in MongoDB using MongoDB Compass, MongoDB Database Tools as well as scripting methods. Additionally, you will have a good idea of how to transfer data between collections and databases.

**Chapter 11: Make It Great - Configuration and Monitoring –** In this chapter, we will consider how the server's start, stop and restart processes work, as well as how to configure your server. We will focus on important settings you will need to know, to properly administrate your server, as well as various monitoring tools and tips for troubleshooting.

**Chapter 12: Seamless Scaling – Replication and Sharding –** In this chapter, we will cover the core concepts of replication in MongoDB, discussing how to setup a replica set and how to leverage the concept of "sharding" for horizontal scaling using MongoDB. By the end of this chapter, you should be comfortable setting up a basic replica set and have a knowledge of how to administrate that replica set. You will also have a solid idea of how sharding works in MongoDB and how and why you would use it for your application.

**Chapter 13: Being Proactive – Security and Backups –** In this chapter, we will begin by discussing how to protect your database via authentication, authorization and roles. Then we will transition to how to effectively backup and restore your databases, wrapping up with a brief discussion about database encryption.

**Chapter 14: Making Stuff – Programming with MongoDB –** By the end of this chapter, you should have a solid idea of how to connect and perform basic queries against MongoDB, using code written in Python, Node.JS (JavaScript) and PHP.

Make sure to see the section about the book's free GitHub Codespace, where you can try some of this code out for yourself, without installing anything locally.

**Chapter 15: Tools for Success – MongoDB Shell and Compass UI –** In this chapter, we will dig a bit deeper into some of the most useful and common tools you will use with MongoDB. We will learn how to configure and personalize the MongoDB Shell, mongosh, as well as how to create useful custom functions you can use within the shell. Then, we will explore the MongoDB Visual Studio Code extension, and MongoDB Playgrounds. Lastly, we will do a bit of a review, as well as a more expansive investigation of MongoDB Compass, the official GUI for MongoDB. By the end of this chapter, you should be empowered to take your use of all these tools to the next level.

**Chapter 16: Cloud Services – MongoDB Atlas –** This chapter will only briefly cover a couple of the key aspects of the cloud services offered by MongoDB Atlas, as indeed, a whole book in itself could be written about Atlas. With that said, after reading this chapter, you should have a solid understanding of what Atlas has to offer and how you can use it to leverage the power of MongoDB even further. Atlas has essentially three main categories of features, two of which we will talk about in this chapter: cloud hosting services and database tools, such as Search and Charts. In the next chapter, we will discuss the third category, Atlas Application Services.

**Chapter 17: MongoDB Atlas – Application Services –** We will be building a React app using Atlas App Services which allows us to build on top of MongoDB, without actually having to run or maintain our own Replica Set, or server. Rather, we will rely on a shared free MongoDB Cluster, Atlas Functions, the Realm SDK for Web and the Atlas' Data API.

**Chapter 18: Jobseeker – Interview Prep –** This chapter will present fifty different interview-like questions, at various levels of difficulty, as well as a response. These may not be the exact questions you would be asked in an interview, for which MongoDB skills are required, but you can expect to be asked some variation of them. For each question, a response is offered, as well as a reference to the chapter we discussed the topic in general, or directly. We will discuss how to use this chapter's questions in the next section.

**Chapter 19: Conclusion –** In this final chapter, we will briefly touch on a few more complex topics such as Change Streams, Transactions in MongoDB, GridFS for storing large files and some comparisons between SQL queries and their MongoDB equivalents.

# Code Bundle and Coloured Images

Please follow the link to download the
***Code Bundle*** and the ***Coloured Images*** of the book:

## https://rebrand.ly/9kqvjnp

The code bundle for the book is also hosted on GitHub at **https://github.com/bpbpublications/MongoDB-For-Jobseekers**. In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at **https://github.com/bpbpublications**. Check them out!

# Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

**errata@bpbonline.com**

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

## Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

## If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

## Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

# Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

**https://discord.bpbonline.com**

# Table of Contents

# Chapter 1

# Why MongoDB?

**"There is a great deal of difference between the eager man who wants to read a book, and the tired man who wants a book to read."**

*-G. K. Chesterton*

## Introduction

Many people love to cook. But do not worry; this is not a book about cooking. We are bringing up food to whet your appetite a little and maybe even illustrate how you are probably using a lot of the concepts we will learn about here, in your everyday lives.

If you are picking up this book, you probably have some idea of what MongoDB is or you might have heard that learning it will help you along in your career. Maybe you just really want to know more: Why should I learn MongoDB; you might be asking?

MongoDB is a powerful, fast, and modern database that was designed for modern applications. From its flexible Document Model (for data), JavaScript based query language, blazing speed, advanced replication and scaling, simple redundancy, advanced aggregation framework, and hooks into the wider MongoDB Atlas Cloud ecosystem, it just might be a clear choice for your next project (or even an existing one)!

Through real life examples and analogies (like cooking), we are going to learn all about MongoDB and its concepts, its associated technologies and how those concepts might apply to the next step in your career.

Hopefully, you are hungry to learn some more. Let us get started.

# Structure

In this chapter, we will discuss the following topics:

- Recipes as data
- The history of data
    - Databases of Data
    - Computer Databases
    - Relational Databases
    - Bringing Separate Data Together
- Data that goes together, can live together
- Thinking of data differently
- This is why MongoDB…

# Objectives

By the end of this chapter, you will have a basic idea of the history of data and databases, how different types of databases work and how MongoDB differs with a newer, modern way to store and interact with data.

# Recipes as Data

We all need food and *good* food that is cooked well, is appreciated throughout the world and across cultures. Food (and cooking it) can sometimes be subjective and complex, yet somehow simple (at its best), at the same time.

What your favourite meals do have in common is some form of structure: there are important details about what ingredients you need (their measurements, units, and so on). There are precise steps to follow, specific temperatures as well as cooking techniques needed, to successfully make the dish. *Figure 1.1* is an example of an old recipe:

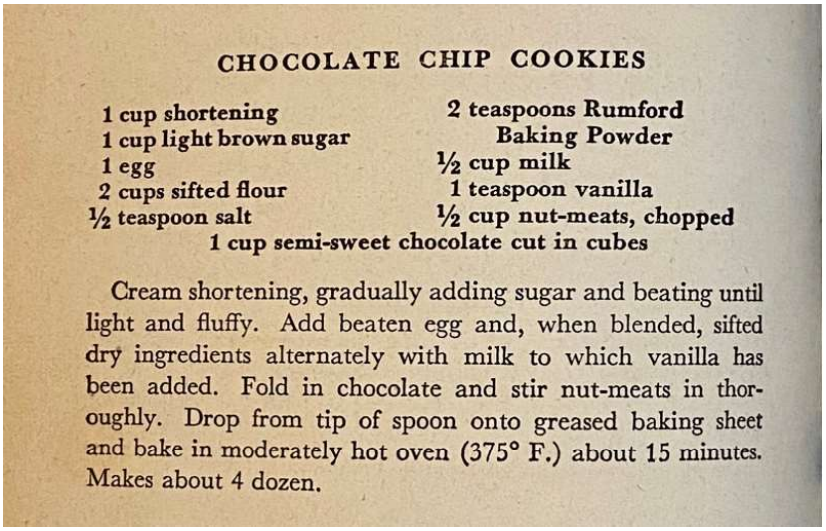

CHOCOLATE CHIP COOKIES

1 cup shortening
1 cup light brown sugar
1 egg
2 cups sifted flour
½ teaspoon salt
2 teaspoons Rumford Baking Powder
½ cup milk
1 teaspoon vanilla
½ cup nut-meats, chopped
1 cup semi-sweet chocolate cut in cubes

Cream shortening, gradually adding sugar and beating until light and fluffy. Add beaten egg and, when blended, sifted dry ingredients alternately with milk to which vanilla has been added. Fold in chocolate and stir nut-meats in thoroughly. Drop from tip of spoon onto greased baking sheet and bake in moderately hot oven (375° F.) about 15 minutes. Makes about 4 dozen.

***Figure 1.1:*** *Century Old Chocolate Chip Cookie Recipe (Public Domain)*

In essence, these instructions are an orderly set of data which we simply call a "recipe." We compile it into a document, which in turn is sometimes gathered into a collection of documents. In simple English, it is called a "Cookbook" or a "Recipe Book".

Switching to a more modern example; what if you wanted a recipe for dinner? You have about half an hour and all you have in your refrigerator, is a chicken. So, you ask the voice assistant on your phone:

**"Find me recipes for chicken that I can cook in 30 minutes or less."**

That might seem like a simple statement but (going with our theme here) there is a lot of information packed into this short sentence. Let us break it out a bit, and see how we might turn this spoken request into a database query:

**"Find me recipes for chicken"**

Starting off, this has the makings of a good query. The user wants us to "find" a particular thing: a "recipe". To make this simple, we will query for recipes literally titled "Chicken". Now, in reality, we would probably do some sort of wildcard search for recipes with the word "chicken" in the title, or more likely also as one of the ingredients.

**"… that I can cook in 30 minutes"**

This is a little more complicated. Do they mean the actual time to cook the recipe is 30 minutes? Or do they mean cooking time plus the preparation time? Given the context here, we will assume that they mean the *combination* of both. Also, this would mean 30 minutes or less, as in less than or equal to.

# Find me a recipe query

Here is how we might construct this as a query in MongoDB:

```
db.recipes.find( { "title" : "Chicken", "total_time" : { $lte: 30 } } );
```

We will gain more information about how querying works and how to compose queries in future chapters. For now, you can see we are looking for a match on the **`title`** field and doing our less than or equal to on a **`total_time`** field. It might look a little funny right now, but the **`$lte`** means "less than or equal to" **`30`** and it is called a Query Operator (they start with a **`$`**). We will learn more about them soon.

Put back into English, it translates to: "Find me recipes where the title is Chicken, and the total time to make it is less than or equal to 30 minutes."

With MongoDB, we will search our recipes, which are stored in **Documents** (more on that in *Chapter 4, A Better Way to Store Data – Documents*) within a **Collection** here called "recipes" (more on that in *Chapter 9, Planning for Performance – Collections and Indexes*) by using the **`find`** method. Given this method of storing data, MongoDB is commonly referred to as a **Document Database**.

Before we go much further, let us step back for a moment and delve a little into the history of data and databases.

# The history of data

**"As a general rule, the most successful man in life is the man who has the best information."**

*-Benjamin Disraeli*

**"Information is the oxygen of the modern age. It seeps through the walls topped by barbed wire; it wafts across the electrified borders."**

*- Ronald Reagan*

There are many types of databases, and the concept of a computer database has been around for almost as long as (electronic) computers themselves. However, data is not a "new" or even "modern" concept. Humans may have only had computers for a couple decades, but data (and even forms of databases) have been around for millennia.

## Databases of clay

The first computer database is generally credited to Charles Bachman in the 1960s while working for IBM. However, the general concept itself goes back thousands