

MySQL for the Internet of Things



Charles Bell

Apress®

MySQL for the Internet of Things

Copyright © 2016 by Charles Bell

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. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

ISBN-13 (pbk): 978-1-4842-1294-3

ISBN-13 (electronic): 978-1-4842-1293-6

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.

Managing Director: Welmoed Spahr

Lead Editor: Jonathan Gennick

Development Editor: Douglas Pundick

Technical Reviewer: Peter Adams

Editorial Board: Steve Anglin, Pramila Balen, Louise Corrigan, Jim DeWolf, Jonathan Gennick,

Robert Hutchinson, Celestin Suresh John, Michelle Lowman, James Markham, Susan McDermott,

Matthew Moodie, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke, Gwenan Spearing

Coordinating Editor: Jill Balzano

Copy Editor: Kim Wimpsett

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

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.springer.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 www.apress.com.

Apress and friends of ED books 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 Special Bulk Sales–eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code/.

*I dedicate this book to my major professors Drs. James E. Ames IV and Lorraine M. Parker
whose tutelage both academic and personal have instilled in me a love for the
art and science of database systems.*

Dr. Charles Bell

Contents at a Glance

About the Author	xiii
About the Technical Reviewer	xv
Acknowledgments	xvii
Introduction	xix
■ Chapter 1: The Internet of Things and Data	1
■ Chapter 2: Hardware for IOT Solutions	29
■ Chapter 3: How IOT Data Is Stored	79
■ Chapter 4: Data Transformation	101
■ Chapter 5: MySQL Primer	141
■ Chapter 6: Building Low-Cost MySQL Data Nodes	195
■ Chapter 7: High Availability IOT Solutions	251
■ Chapter 8: Demonstration of High Availability Techniques	273
Index	311

Contents

About the Author	xiii
About the Technical Reviewer	xv
Acknowledgments	xvii
Introduction	xix
■ Chapter 1: The Internet of Things and Data	1
IOT Solutions	2
IOT Is More Than Just Connected to the Internet.....	3
IOT Services.....	6
Example IOT Solutions	7
What Is IOT Data?	11
IOT Predictions: Data Overload?	16
Addressing IOT Devices	16
IOT and Big Data	19
IOT Security	20
Common Security Threats	23
Securing IOT Solutions	25
Summary	28
■ Chapter 2: Hardware for IOT Solutions	29
Microcontrollers	29
What Is an Arduino?.....	30
Arduino Models.....	31
Arduino Clones	36
Arduino Tutorial.....	45
Additional Arduino Hardware	52

Low-Powered Computing Platforms	62
Arduino Hybrids	63
Computer Boards	66
Sensors	72
Analog Sensors	73
Digital Sensors	73
Storing Sensor Data	74
Examples of Sensors	74
Computer Systems	77
Summary	78
■ Chapter 3: How IOT Data Is Stored	79
Distributed IOT	80
Data Collectors	81
Data Collectors with Storage	81
Actionable Device	81
Data Aggregators	81
Database Server	82
Local On-Device Storage	82
Local Storage on the Raspberry Pi	83
Local Storage on the Arduino	85
Passing the Buck to Aggregators	90
Database Storage	92
Benefits	93
Techniques	94
Considerations	95
Distributed IOT Network Best Practices	96
Node Placement	96
Data Storage	97
Presentation	99
Summary	99

■ Chapter 4: Data Transformation	101
Making Sense of IOT Data	102
What Is Being Observed?.....	102
Is There Another Way to Make the Observation?	102
How Often Do You Need to Record the Observation?	103
What Type of Data Does the Sensor Produce?.....	104
Are There Interpretations Needed for the Observation Data?.....	104
What Level of Accuracy Do You Need?	105
What Is the Lifetime of the Data?	105
Annotation	106
Recording the Sensor Name or Adding Notes	107
Recording the Date and Time	110
Data Type Transformations	116
Adding Derived or Calculated Data.....	123
Data Interpretations.....	130
Aggregation.....	134
Data from Multiple Sensors.....	134
Data from Multiple Nodes.....	137
Aggregate Calculations.....	139
Summary.....	140
■ Chapter 5: MySQL Primer	141
Getting Started	141
How Do I Use MySQL?	142
How to Get and Install MySQL	144
How Data Is Stored and Retrieved.....	157
How and Where MySQL Stores Data.....	158
Common MySQL Commands and Concepts	164
MySQL Commands	164
MySQL Concepts.....	171

Planning Database Storage for IOT Data	176
Example 1: Plant-Monitoring System	176
Recommendations and Best Practices	191
Summary	193
■ Chapter 6: Building Low-Cost MySQL Data Nodes	195
Introducing the Raspberry Pi	195
Noble Origins	197
Models	198
A Tour of the Board	200
Required Accessories	201
Recommended Accessories	202
Where to Buy	203
Raspberry Pi Tutorial	206
Choosing a Boot Image (Operating System)	206
Booting Up	212
MySQL Installation and Setup	215
Partitioning and Formatting the Drive	216
Setting Up Automatic Drive Mounting	218
Installing MySQL Server	220
Other Platforms	225
BeagleBone Black	225
pcDuino	226
Intel Galileo	227
MySQL Clients: How to Connect and Save Data	230
Introducing Connector/Arduino	230
Introducing Connector/Python	242
Summary	249

■ Chapter 7: High Availability IOT Solutions	251
What Is High Availability?	251
High Availability Options for IOT Solutions with MySQL	253
Recovery	253
Redundancy	254
Scaling	255
Fault Tolerance	255
High Availability Techniques	256
Backup and Recovery	256
MySQL Replication Primer	264
Fault Tolerance in IOT Nodes	270
Summary	271
■ Chapter 8: Demonstration of High Availability Techniques.....	273
MySQL Replication Techniques	273
Transaction Processing.....	274
Advanced Replication with Global Transaction Identifiers.....	276
Replication and Database Maintenance Tips.....	282
Example: Scaling Applications.....	285
High Availability IOT Nodes.....	291
Example: Redundant Data Collectors.....	291
Example: Fault-Tolerant Data Collector	300
Summary	310
Index.....	311

About the Author



Dr. Charles Bell conducts research in emerging technologies. He is a member of the Oracle MySQL Development team as a development manager directing development of MySQL high availability solutions. He lives in a small town in rural Virginia with his loving wife. He received his doctorate of philosophy in engineering from Virginia Commonwealth University in 2005. His research interests include database systems, software engineering, sensor networks, and 3D printing. He spends his limited free time as a practicing Maker focusing on microcontroller and 3D printers and printing projects.

About the Technical Reviewer



Peter Adams has been crafting web and mobile applications for businesses of every size since 1998. To make a living, he manages a small development team working with businesses to build their technology stack. To live when he's not behind the keyboard, he takes advantage of the Arizona climate to hike, mountain bike, and motorcycle year round. Occasionally, he takes time to write about his adventures at www.mymegaverse.org/blog. More often, you can find him keeping his clients updated at www.pingdevelopment.com.

Acknowledgments

I would like to thank all of the many talented and energetic professionals at Apress. I appreciate the understanding and patience of my editor, Jonathan Gennick, and managing editor, Jill Balzano. They were instrumental in the success of this project. I would also like to thank the army of publishing professionals at Apress for making me look so good in print. Thank you all very much!

I'd like to especially thank the technical reviewer, Peter Adams, for his often-profound insights, constructive criticism, and encouragement. I'd also like to thank my friends for their encouragement and suggestions for things to include in the book.

Most importantly, I want to thank my wife, Annette, for her unending patience and understanding while I spent so much time with my laptop.

Introduction

Internet of Things (IoT) solutions are not nearly as complicated as the name may seem to indicate. Indeed, the IoT is largely another name for what we have already been doing. You may have heard of “connected devices” or “Internet-ready” or even “cloud-enabled.” All of these refer to the same thing—be it a single device such as a toaster or a plant monitor or a complex, multidevice product like home automation solutions. They all share one thing in common: they can be accessed via the Internet to either display data or interact with the devices directly. The trick is applying knowledge of technologies to leverage them to the best advantages for your IoT solution. In this book, we explore how to leverage MySQL in your IoT solution as a means to store your data.

Intended Audience

I wrote this book to share my passion for MySQL and IoT solutions. I especially wanted to show how anyone can use MySQL to store IoT data—even from small microcontrollers. Indeed, I want to share with the world the power and ease of installation and use of MySQL, the world’s most popular open source database.

The intended audience therefore includes anyone interested in learning how to install and use MySQL, hobbyists and enthusiasts who need a way to store data in their IoT solutions, and designers and engineers building commercial IoT solutions.

How This Book Is Structured

The book was written to guide the reader from a general knowledge of MySQL to expertise in developing MySQL solutions for the IoT. The first several chapters cover general topics including a short introduction to the Internet of Things, how data is generated and stored, and some of the available hardware for IoT. Additional chapters present a primer on MySQL including how to install and configure database servers. Throughout the book are examples of how to implement many of the concepts presented. The following is a brief overview of each chapter included in this book.

- **Chapter 1, “The Internet of Things and Data”:** This chapter presents and answers the questions of what the IOT is and how IOT solutions are constructed. You are introduced to some terminology to describe the architecture of IOT solutions as well as some examples of well-known IOT solutions. The chapter concludes with a discussion of the two most critical issues of IOT solutions through practical examples: big data and security.

- **Chapter 2, “Hardware for IOT Solutions”:** This chapter discusses some of the hardware available for building IoT solutions including the Arduino line of microcontroller boards, low-cost (low-power) computing boards such as the Raspberry Pi, and the various communication hardware you can use to connect the nodes from using an Ethernet or WiFi network connecting to the Internet to low-cost, low-power wireless communication modules (XBee radios) for connecting sensor nodes. The chapter also presents a tutorial of the Arduino programming environment along with examples of how to use the Arduino with sensors. The chapter concludes with a brief look at the types of sensors available for building IoT solutions.
- **Chapter 3, “How IOT Data Is Stored”:** This chapter examines some of the methods available to you for storing data. You will see examples of how to read and write data in files on the Arduino and Raspberry Pi. The chapter also presents the benefits, considerations, and recommendations for deploying a database server in your IOT solution. The chapter concludes with a list of the best practices for designing a network of nodes for your IOT solution.
- **Chapter 4, “Data Transformation”:** This chapter presents practical questions to ask when considering the data as well as several examples of annotation and aggregation in both Arduino and Python code. This chapter also includes considerations for storing IOT data in a database such as implementing annotations and aggregations in the database rather than in code.
- **Chapter 5, “MySQL Primer”:** This chapter dives into discovering the power of using a database server, learning how database servers store data, and seeing how to issue commands for creating databases and tables for storing data as well as commands for retrieving that data. The chapter is a primer on MySQL including how to get started with your own IOT data through database and code examples.
- **Chapter 6, “Building Low-Cost MySQL Data Nodes”:** This chapter presents a crash course on how to set up a Raspberry Pi, install MySQL, and use them. The chapter also includes examples on how to write data to a MySQL database server using an Arduino sketch and a Python program on another machine (Raspberry Pi, BeagleBone Black, pcDuino, and so on). The chapter concludes with a brief look at the high availability features of MySQL.
- **Chapter 7, “High Availability IOT Solutions”:** This chapter discusses what high availability is and how high availability concepts can be realized. You will also learn key high availability concepts, tools, and techniques for MySQL including backup and recovery and replication. You also will discover how to implement fault tolerance for collecting data on microcontrollers.
- **Chapter 8, “Demonstration of High Availability Techniques”:** This chapter presents more about MySQL replication including some helpful tips and techniques for setting up and using replication in your IOT solutions. The chapter also presents some examples of high availability concepts including a simple round-robin read scale-out solution for Python, hardware for building a redundant data collector with failover, and how to build in fault tolerance for data collectors so that you don’t lose data should the communication pathway from data collector to database fail.

How to Use This Book

This book is designed to guide you through learning more about what the Internet of Things is, discovering the power of MySQL and database systems, and seeing how to build your IoT solutions so that they are more reliable by leveraging high availability techniques.

If you are familiar with some of the topics early in the book, I recommend you skim them so that you are familiar with the context presented so that the later chapters, especially the examples, are easy to understand and implement on your own. You may also want to read some of the chapters out of order so that you can get your project moving, but I recommend going back to the chapters you skip to ensure you get all of the data presented.

If you are just getting started with MySQL and your IoT solutions or are not well versed in database systems, I recommend reading the book in its entirety before developing your own IoT solution. That said, many of the examples permit you to build small examples that you can use to help learn the concepts. For example, you can learn how to read sensors, store data, build a data aggregator, and more, for sensor networks and IoT solutions.

Downloading the Code

The code for the examples shown in this book is available on the Apress web site, www.apress.com. You can find a link on the book's information page on the Source Code/Downloads tab. This tab is located in the Related Titles section of the page.

Contacting the Author

Should you have any questions or comments—or even spot a mistake you think I should know about—you can contact the author at drcharlesbell@gmail.com.