

THE EXPERT'S VOICE® IN DATABASES

SECOND EDITION

Expert MySQL

APPLY MASTERY OF MySQL INTERNALS
TO HIGH-AVAILABILITY, REPLICATION,
CUSTOMIZATION, AND BEYOND

Charles Bell

Apress®

Expert MySQL

Second Edition



Charles Bell

Apress®

Expert MySQL: Second Edition

Copyright © 2012 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 978-1-4302-4659-6

ISBN 978-1-4302-4660-2 (eBook)

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 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. The publisher makes no warranty, expressed or implied, with respect to the material contained herein.

President and Publisher: Paul Manning

Lead Editor: Jonathan Gennick

Technical Reviewers: Alfranio Correia, Sven Sandberg, Luis Soares, Marco Tusa, Geert Vanderkelen

Editorial Board: Steve Anglin, Ewan Buckingham, Gary Cornell, Louise Corrigan, Morgan Ertel, Jonathan

Gennick, Jonathan Hassell, Robert Hutchinson, Michelle Lowman, James Markham, Matthew Moodie,

Jeff Olson, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke, Dominic Shakeshaft, Gwenan Spearing,

Matt Wade, Tom Welsh

Coordinating Editor: Kevin Shea

Copy Editor: Pat Morris

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Cover Designer: Anna Ishchenko

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.

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 material 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 biggest fan—my big sister.
You were the catalyst for my academic and writing success.*

—Charles Bell

Contents at a Glance

About the Author	xix
About the Technical Reviewers	xxi
Acknowledgments	xxiii
Introduction	xxv
■ Part 1: Getting Started with MySQL Development	1
■ Chapter 1: MySQL and The Open Source Revolution	3
■ Chapter 2: The Anatomy of a Database System	23
■ Chapter 3: A Tour of the MySQL Source Code.....	57
■ Chapter 4: Test-Driven MySQL Development	117
■ Part 2: Extending MySQL	151
■ Chapter 5: Debugging	153
■ Chapter 6: Embedded MySQL	195
■ Chapter 7: Adding Functions and Commands to MySQL	251
■ Chapter 8: Extending MySQL High Availability	281
■ Chapter 9: Developing MySQL Plugins.....	339
■ Chapter 10: Building Your Own Storage Engine.....	369
■ Part 3: Advanced Database Internals	453
■ Chapter 11: Database System Internals	455
■ Chapter 12: Internal Query Representation	465

■ CONTENTS AT A GLANCE

■ Chapter 13: Query Optimization	495
■ Chapter 14: Query Execution	543
■ Appendix.....	587
Index.....	601

Contents

About the Author	xix
About the Technical Reviewers	xxi
Acknowledgments	xxiii
Introduction	xxv
■ Part 1: Getting Started with MySQL Development	1
■ Chapter 1: MySQL and The Open Source Revolution	3
What Is Open Source Software?.....	4
Why Use Open Source Software?	5
Is Open Source Really a Threat to Commercial Software?	8
Legal Issues and the GNU Manifesto	9
Let the Revolution Continue!	11
Developing with MySQL	12
Why Modify MySQL?.....	14
What Can You Modify in MySQL? Are There Limits?	15
MySQL's Dual License.....	16
So, Can You Modify MySQL or Not?	18
Guidelines for Modifying MySQL.....	18
A Real-World Example: TiVo	19
Summary	20

■ Chapter 2: The Anatomy of a Database System	23
Types of Database Systems	23
Object-Oriented Database Systems.....	23
Object-Relational Database Systems	24
Relational Database Systems.....	25
Relational Database System Architecture	27
Client Applications	27
Query Interface	29
Query Processing.....	30
Query Optimizer.....	32
Internal Representation of Queries.....	34
Query Execution.....	35
File Access.....	36
Query Results	38
Relational Database Architecture Summary.....	38
The MySQL Database System	38
MySQL System Architecture	39
SQL Interface	41
Parser	41
Query Optimizer.....	42
Query Execution.....	43
Query Cache	43
Cache and Buffers	46
File Access via Pluggable Storage Engines	49
Summary.....	56
■ Chapter 3: A Tour of the MySQL Source Code.....	57
Getting Started	57
Understanding the Licensing Options.....	57
Getting the Source Code.....	58

The MySQL Source Code	59
Getting Started	60
The mysqld_main() Function	62
Handling Connections and Creating Threads.....	64
Parsing the Query	71
Preparing the Query for Optimization.....	79
Optimizing the Query	83
Executing the Query	88
Supporting Libraries	91
Important Classes and Structures	92
MySQL Plugins.....	99
Coding Guidelines.....	102
General Guidelines.....	102
Documentation	102
Functions and Parameters.....	104
Naming Conventions.....	105
Spacing and Indenting.....	105
Documentation Utilities	106
Keeping an Engineering Logbook	108
Tracking Your Changes	109
Building the System for the First Time	111
Summary	115
■ Chapter 4: Test-Driven MySQL Development	117
Background	117
Why Test?	117
Benchmarking	120
Profiling	122
Introducing Software Testing.....	123
Functional Testing vs. Defect Testing	123

MySQL Testing	127
Using the MySQL Test Suite	128
MySQL Benchmarking	136
MySQL Profiling	143
Summary	149
Part 2: Extending MySQL	151
Chapter 5: Debugging	153
Debugging Explained	153
The origins of debugging	154
Debugging Techniques	154
Basic Process	154
Approaches to Debugging	156
Inline Debugging Statements	156
Error Handlers	159
External Debuggers	161
Stand-alone Debuggers	161
Interactive Debuggers	164
GNU Data Display Debugger	166
Bidirectional Debuggers	168
Debugging MySQL	170
Inline Debugging Statements	170
Error Handlers	178
Debugging in Linux	179
Using gdb	179
Using ddd	183
Debugging in Windows	189
Summary	193
Chapter 6: Embedded MySQL	195
Building Embedded Applications	195
What Is an Embedded System?	195

Types of Embedded Systems	196
Embedded Database Systems.....	196
Embedding MySQL	197
Methods of Embedding MySQL	198
Bundled Server Embedding.....	198
Deep Embedding (<code>libmysqld</code>).....	198
Resource Requirements	198
Security Concerns	199
Advantages of MySQL Embedding	199
Limitations of MySQL Embedding	199
The MySQL C API	200
Getting Started	200
Most Commonly Used Functions	201
Creating an Embedded Server	202
Initializing the Server	203
Setting Options.....	204
Connecting to the Server.....	205
Running Queries.....	206
Retrieving Results	206
Cleanup	207
Disconnecting from and Finalizing the Server	208
Putting It All Together	208
Error Handling	209
Building Embedded MySQL Applications.....	210
Compiling the Library (<code>libmysqld</code>)	210
Compiling <code>libmysqld</code> on Linux.....	210
Compiling <code>libmysqld</code> on Windows.....	210
What About Debugging?	211
What About the Data?.....	213

Creating a Basic Embedded Server.....	213
Linux Example	213
Windows Example	217
What About Error Handling?	223
Embedded Server Application	224
The Interface	224
The Data and Database	227
Creating the Project.....	228
Design.....	230
Managed vs. unmanaged code	230
Database Engine Class	230
Customer Interface (Main Form).....	238
Administration Interface (Administration Form)	244
Detecting Interface Requests	247
Compiling and Running	247
Summary.....	249
■ Chapter 7: Adding Functions and Commands to MySQL	251
 Adding User-Defined Functions.....	251
CREATE FUNCTION Syntax	252
DROP FUNCTION Syntax	252
Creating a User-Defined Library	252
Adding a New User-Defined Function.....	257
 Adding Native Functions	262
 Compiling and Testing the New Native Function	266
 Adding SQL Commands.....	267
 Adding to the Information Schema.....	275
 Summary.....	280

Chapter 8: Extending MySQL High Availability	281
What is Replication?.....	281
Why use Replication?	282
How Does Replication Achieve High Availability?	283
Basic Replication Setup	284
Requirements for Replication	284
Configuring the Master.....	286
Configuring the Slave	288
Connecting the Slave to the Master	288
Next Steps	290
The Binary Log	290
Row Formats	291
The mysqlbinlog Client.....	291
SHOW BINLOG EVENTS Command	293
Additional Resources.....	296
Replication Architecture	296
A Brief Tour of the Replication Source Code.....	297
Replication Source Code Files	297
Log Events Explained.....	299
Types of Log Events.....	305
Execution of Log Events	307
Extending Replication.....	310
Global Slave Stop Command	311
Slave Connect Logging	322
Summary	338
Chapter 9: Developing MySQL Plugins.....	339
MySQL Plugins Explained	339
Types of Plugins.....	339
Using MySQL Plugins.....	340
The MySQL Plugin API.....	344
Compiling MySQL Plugins.....	347

The RFID Authentication Plugin	348
Concept of Operations	349
RFID Module	349
Architecture of Authentication Plugins	355
Building the RFID Authentication Plugin	357
RFID Authentication in Action	365
Further Work.....	367
Summary.....	368
■ Chapter 10: Building Your Own Storage Engine.....	369
MySQL Storage Engine Overview	369
Storage Engine Development Process	371
Source Files Needed	372
Unexpected Help	372
The Handlerton.....	372
The Handler Class	375
A Brief Tour of a MySQL Storage Engine	379
The Spartan Storage Engine.....	380
Low-Level I/O Classes	380
The Spartan_data Class.....	381
The Spartan_index Class	389
Getting Started	403
Stage 1: Stubbing the Engine	405
Creating the Spartan Plugin Source Files.....	405
Adding the CMakeLists.txt File	406
Compiling the Spartan Engine	407
Testing Stage 1 of the Spartan Engine	407
Stage 2: Working with Tables	411
Updating the Spartan Source Files	412
Updating the Class File	413

Testing Stage 2 of the Spartan Engine	416
Stage 3: Reading and Writing Data.....	417
Updating the Spartan Source Files	417
Testing Stage 3 of the Spartan Engine	421
Stage 4: Updating and Deleting Data	422
Updating the Spartan Source Files	423
Testing Stage 4 of the Spartan Engine	424
Stage 5: Indexing the Data	429
Updating the Spartan Source Files.....	430
Testing Stage 5 of the Spartan Engine	442
Stage 6: Adding Transaction Support	448
Summary.....	452
Part 3: Advanced Database Internals	453
Chapter 11: Database System Internals	455
Query Execution	455
MySQL Query Execution Revisited.....	455
What Is a Compiled Query?	456
Exploring MySQL Internals	457
Getting Started Using MySQL for Experiments	457
Limitations and Concerns.....	459
The Database System Internals Experiment	460
Why an Experiment?.....	460
Overview of the Experiment Project	460
Components of the Experiment Project	461
Conducting the Experiments.....	463
Summary.....	463