

《高性能MySQL》

图书基本信息

书名：《高性能MySQL》

13位ISBN编号：9787564134457

10位ISBN编号：7564134453

出版时间：2012-6

出版社：东南大学出版社

页数：793

版权说明：本站所提供下载的PDF图书仅提供预览和简介以及在线试读，请支持正版图书。

更多资源请访问：www.tushu000.com

前言

We wrote this book to serve the needs of not just the MySQL application developer but also the MySQL database administrator. We assume that you are already relatively experienced with MySQL. We also assume some experience with general system administration, networking, and Unix-like operating systems. The second edition of this book presented a lot of information to readers, but no book can provide complete coverage of a topic. Between the second and third editions, we took notes, on literally thousands of interesting problems we'd solved or seen others solve. When we started to outline the third edition, it became clear that not only would full coverage of these topics require three to five thousand pages, but the book still wouldn't be complete. After reflecting on this problem, we realized that the second edition's emphasis on deep coverage was actually self-limiting, in the sense that it often didn't teach readers how to think about MySQL. As a result, this third edition has a different focus from the second edition. We still convey a lot of information, and we still emphasize the same goals, such as reliability and correctness. But we've also tried to imbue the book with a deeper purpose: we want to teach the principles of why MySQL works as it does, not just the facts about how it works. We've included more illustrative stories and case studies, which demonstrate the principles in action. We build on these to try to answer questions such as "Given MySQL's internal architecture and operation, what practical effects arise in real usage? Why do those effects matter? How do they make MySQL well suited (or not well suited) for particular needs?" Ultimately, we hope that your knowledge of MySQL's internals will help you in situations beyond the scope of this book. And we hope that your newfound insight will help you to learn and practice a methodical approach to designing, maintaining, and troubleshooting systems that are built on MySQL.

How This Book Is Organized

We fit a lot of complicated topics into this book. Here, we explain how we put them together in an order that makes them easier to learn.

A Broad Overview

Chapter 1, MySQL Architecture and History is dedicated to the basics--things you'll need to be familiar with before you dig in deeply. You need to understand how MySQL is organized before you'll be able to use it effectively. This chapter explains MySQL's architecture and key facts about its storage engines. It helps you get up to speed if you aren't familiar with some of the fundamentals of a relational database, including transactions. This chapter will also be useful if this book is your introduction to MySQL but you're already familiar with another database, such as Oracle. We also include a bit of historical context: the changes to MySQL over time, recent ownership changes, and where we think it's headed.

Building a Solid Foundation

The early chapters cover material we hope you'll reference over and over as you use MySQL.

Chapter 2, Benchmarking MySQL

discusses the basics of benchmarking--that is, determining what sort of workload your server can handle, how fast it can perform certain tasks, and so on. Benchmarking is an essential skill for evaluating how the server behaves under load, but it's also important to know when it's not useful.

Chapter 3, Profiling Server Performance

introduces you to the response time-oriented approach we take to troubleshooting and diagnosing server performance problems. This framework has proven essential to solving some of the most puzzling cases we've seen. Although you might choose to modify our approach (we developed it by modifying Cary Millsap's approach, after all), we hope you'll avoid the pitfalls of not having any method at all.

In Chapters 4 through 6, we introduce three topics that together form the foundation for a good logical and physical database design. In Chapter 4, Optimizing Schema and Data Types, we cover the various nuances of data types and table design. Chapter 5, Indexing for High Performance extends the discussion to indexes--that is, physical database design. A firm understanding of indexes and how to use them well is essential for using MySQL effectively, so you'll probably find yourself returning to this chapter repeatedly. And Chapter 6, Query Performance Optimization wraps the topics together by explaining how MySQL executes queries and how you can take advantage of its query optimizer's strengths. This chapter also presents specific examples of many Common classes of queries, illustrating where MySQL does a good job and how to transform queries into forms that use its strengths.

Up to this point, we've covered the basic topics that apply to any database: tables, indexes, data, and queries. Chapter 7, Advanced MySQL Features goes beyond the basics and shows you how MySQL's advanced features work. We examine topics such as partitioning, stored procedures, triggers, and character sets. MySQL's implementation of these features is different from other databases, and a good understanding of them can open up new opportunities for performance gains that you might not have thought about otherwise.

Configuring Your

Application The next two chapters discuss how to make MySQL, your application, and your hardware work well together. In Chapter 8, *Optimizing Server Settings*, we discuss how you can configure MySQL to make the most of your hardware and to be reliable and robust. Chapter 9, *Operating System and Hardware Optimization* explains how to get the most out of your operating system and hardware. We discuss solid-state storage in depth, and we suggest hardware configurations that might provide better performance for larger-scale applications.

Both chapters explore MySQL internals to some degree. This is a recurring theme that continues all the way through the appendixes: learn how it works internally, and you'll be empowered to understand and reason about the consequences: MySQL as an Infrastructure Component MySQL doesn't exist in a vacuum. It's part of an overall application stack, and you'll need to build a robust overall architecture for your application. The next set of chapters is about how to do that. In Chapter 10, *Replication*, we discuss MySQL's killer feature: the ability to set up multiple servers that all stay in sync with a master server's changes. Unfortunately, replication is perhaps MySQL's most troublesome feature for some people. This doesn't have to be the case, and we show you how to ensure that it keeps running well. Chapter 11, *Scaling MySQL* discusses what scalability is (it's not the same thing as performance), why applications and systems don't scale, and what to do about it. If you do it right, you can scale MySQL to suit nearly any purpose. Chapter 12, *High Availability* delves into a related-but-distinct topic: how to ensure that MySQL stays up and functions smoothly. In Chapter 13, *MySQL in the Cloud*, you'll learn about what's different when you run MySQL in cloud computing environments.

《高性能MySQL》

内容概要

《高性能MySQL(影印版)(第3版)(英文版)》学习MySQL5.5版提供的新特性，包括存储过程、数据库分区、触发器和视图，实现在复制、高可用性和集群上的改进，实现MySQL在云环境中运行的高性能，优化高级查询特性，如全文检索，充分利用多核处理器和固态硬盘等硬件带来的优势，探索备份和恢复策略——包括新的在线热备份工具。

作者简介

作者：（美国）施瓦兹（Baron Schwartz）（美国）扎伊采夫（Peter Zaitsev）（美国）特卡申科（Vadim Tkachenko）施瓦兹（Schwartz B.），是Percona公司的首席性能架构师。他的主要工作是通过创建一系列工具和技术来提高MySQL数据库的易用性和可靠性。扎伊采夫（zaitsev P.），是Percona公司的首席执行官和联合创始人之一。作为一名数据库内核、计算机硬件和应用扩展技术领域的专家，他管理着MySQL团队内部的高性能技术组直到2006年。特卡申科（Tkachenko V.）是Percona公司的首席技术官和联合创始人之一。他领导着公司的开发团队先后开发了Percona服务器、Percona XtraDB集群和Percona XtraBackup工具等产品。

书籍目录

Foreword

Preface

1. MySQL Architecture and History

2. Benchmarking MySQL

3. Profiling Server Performance

4. Optimizing Schematic and Data Types

5. Indexing for High Performance

6. Query Performance Optimization

7. Advanced MySQL Features

8. Optimizing Server Setting

9. Operating System and Hardware Optimization

10. Replication

11. Scaling MySQL

12. High Availability

13. MySQL in the Cloud

14. Application-Level Optimization

15. Backup and Recovery

16. Tools for MySQL Use

A. Forks and Variants of MySQL

B. MySQL Server Status

C. Trafferring Large Files

D. Using EXPLAIN

章节摘录

版权页：插图： Too many columns MySQL's storage engine API works by copying rows between the server and the storage engine in a row buffer format;the server then decodes the buffer into col-umns.But it can be costly to turn the row buffer into the row data structure with the decoded columns.MyISAM's fixed row format actually matches the server's row format exactly,so no conversion is needed.However,MyISAM's variable row format and InnoDB's row format always require conversion.The cost of this con-version depends on the number of columns.We discovered that this can become expensive when we investigated an issue with high CPU consumption for a cus-tomer with extremely wide tables (hundreds of columns),even though only a few columns were actually used.If you're planning for hundreds of columns,be aware that the server's performance characteristics will be a bit different. Too many joins The so-called entity-attribute-value (EAV) design pattern is a classic case of a uni-versally bad design pattern that especially doesn't work well in MySQL.MySQL has a limitation of 61 tables per join,and EAV databases require many self-joins. We've seen more than a few EAV databases eventually exceed this limit.Even at many fewer joins than 61,however,the cost of planning and optimizing the query can become problematic for MySQL.As a rough rule of thumb,it's better to have a dozen or fewer tables per query if you need queries to execute very fast with high concurrency. The all-powerful ENUM Beware of overusing ENUM.Here's an example we saw : `CREATE TABLE ... (country enum('','0','1','2',..., '31')` The schema was sprinkled liberally with this pattern.This would probably be a questionable design decision in any database with an enumerated value type,be-cause it really should be an integer that is foreign-keyed to a "dictionary" or "lookup" table anyway.But in MySQL,you can't add a new country to the list without an `ALTER TABLE`,which is a blocking operation in MySQL 5.0 and earlier,and even in 5.1 and newer if you add the value anywhere but at the end of the list.

《高性能MySQL》

媒体关注与评论

“同本书的前几版相比，第三版是一本更加棒的书。这几位作者是唯一具有资格写这本书的人。我将继续向他们学习，并希望你能抽出时间这样做。” ——Mark Callaghan Facebook软件工程师

精彩短评

- 1、如果你有英语基础的话，就选择这本书吧，你可以得到很多。目前看了几章，之后再追回复。
- 2、新版加入了云计算相关的内容，虽然是英文原版，但看着也不是很累。只是书籍的纸张，感觉有点粗糙了。
- 3、刚看了前几张，收获很多，得经常翻阅。
- 4、经典书籍还是要看影音版的，值得收藏
- 5、挺喜欢这本书，书本身印刷质量不错，送货速度不错。
- 6、正在学习中，好厚的一本书，质量感觉没有那么好
- 7、毕竟是影印版的，要求也不高，自己清晰，没错就行。中文版的大致看了一遍，感觉有些翻译的不好。看不明白，所以买本原本的比对着看。这本书可算得上是MYSQL的经典书籍啊，不过对于我这种初学者，有些部分还是不适用的。只能等以后慢慢体会了。
- 8、书很好，很受用！也锻炼了英文。
- 9、还是比较简单易懂的，就是云计算那章有点晕，一些名词挺纠结的。。
- 10、居然买了本英文版书！
- 11、书的内容很好，学到很多东西
- 12、不错,了解mysql的好书
- 13、全是英文，看上去有中麻辣辣的感觉呀~很厚的一本书，对于学计算机的来说很不错喜欢~里面的纸张一般吧，但是搞活动还是蛮划算的！
- 14、好书，就是英文版，不过大学生读无压力，建议学数据库的人都读一读
- 15、书挺好，值得仔细研究。
- 16、09年9月，我买了这本书的第二版，也是英文版，爱不释手，希望这第三版会更好~~
- 17、还没怎么看，书很不错
- 18、之前看过电子版的中文的本想买本书做手册使结果买了个英文版的看着好吃力哈哈内容还是很不错的
- 19、书蛮好的~~
- 20、super.
- 21、这个是英文的影印版，本来想弄本中文的看看的。。
- 22、用mysql必读。写的清晰而专业。
- 23、详细的介绍了Mysql的性能测试、内部机制、SQL优化、Index、Master Slave部署
- 24、朋友说不错,值得看一下
- 25、比较懒，还差3章没有看完，收获很大
- 26、英文版，英文不好的可以等翻译版出来
- 27、看过了，还可以，值得好好学习一下
- 28、但是纸张偏暗 外表红色有点山寨。还是没有机械工业出版社的影印版来的讲究。但是只有这个出版社现在有这款书。还算出的快 这个价钱，我还是很推荐的
- 29、纸质一般,内容应该不错

版权说明

本站所提供下载的PDF图书仅提供预览和简介，请支持正版图书。

更多资源请访问:www.tushu000.com