

# Welcome

Call this stage three of the computer revolution. Stage one worked out the basic principles of computing, and applied them to mainframes that helped run the world. That was pretty much in hand by 1980. Stage two turned those principles into software you could install on your desktop or laptop for a reasonable price. That was in hand by 1995. We are in stage three, which networks billions of desktops, laptops and mobiles, supplying them with *free open source software development tools for any computer task you can imagine*.

MySQL<sup>®</sup> is one of those tools. It runs a majority of the world's database-driven websites. On a suite of world popularity measures it *ranks* 2, just behind Oracle and just ahead of SQL Server. Version 5.0 came into production in 2005 and has been archived, 5.1 came in 2008, 5.5 in late 2011, 5.6 in early 2013, 5.7 in late 2015. The next major version, 8.0, has seen its first development milestone release. MySQL is fast, reliable, flexible, richly featured and customisable. It's built for developers, it's robust, and it's open source. Under a GPL/GNU licence it costs you exactly nothing.

MySQL queries typically return results in less than 0.05 seconds. Well-designed databases up to 1.5 terabytes perform well without custom tweaks like sharding. Replication is straightforward. The relational InnoDB storage engine is robust. So is the server overall: Google engineer Jeremy Cole found a Twitter MySQL installation that had been running 212 days without downtime, over which time it had processed an average of 6.9 queries per second, returning about 1.4 million data rows per second.

Fine. But if you are a database developer, you probably already know about MySQL. Why another book about it?

Three answers: about recent MySQL versions, about our approach, and about you.

*About recent versions:* In 5.0 and 5.1, correlated subqueries, stored routines including Triggers, Views, and an `information_schema` implementation took MySQL well beyond the one-app-one-database model of early MySQL versions. Versions 5.5 and 5.6 made transaction management more robust, added tools, and improved performance as much as 200%. 5.7 brings better security, logging and Triggers; subqueries in Views; JSON and derived columns, and many bug fixes. If Oracle's limited commitment to open source principles worries you, Monty Widenius's MariaDB has drop-in replacement versions for MySQL 5.5 and 5.6.

*About our approach:* MySQL's original focus was on a small footprint and high performance, sometimes at the cost of relational database correctness. MySQL databases tended to be maintained from single web sites. You can't argue with what MySQL AB accomplished, but more MySQL users began to ask for big-three-style relational power. In short order MySQL AB declared it was aiming at SQL:2003 compatibility, was bought

by a big software company (Sun), which in 2009 was itself bought by a big three database company, Oracle.

So in this book, instead of beginning with how to do up a limited MySQL project as quickly as possible, we ground design on use case analysis, we ground syntax on relational database theory, and we ground database maintenance on software life cycle analysis. Most everything else follows from those fundamentals. You spend a bit more time up front on the basics, and soon you find you can drive the vehicle a lot faster, and more steadily, to more interesting destinations.

*About you:* You may be new to MySQL. You may be a Windows/SQL Server developer or manager who now, for the first time, needs to go to Linux or MySQL.

You may have decided that for you or your working community, the time for database systems from huge corporations costing tens of thousands of dollars has passed, that the time for community-developed open-source software development has arrived.

You may work for an NGO with a certain amount of in-house skill and an insistence on being honest, but now lack the funds for expensive licences. You may have worked mostly on Oracle or SQL Server or Sybase or even FoxPro systems, but your next project specifies MySQL. Whatever the reason, you are, or soon will be using a major open-source RDBMS *which is changing as you use it, on demand from you.*

Traditional dead-tree publishing offers a 500-1000 page tome of which maybe 200 pages are relevant to you. Some of it is out of date by the time you read it. So you lug home this two-kilogram object, read the pages of interest, curse the out-of-dateness, and ignore the rest. Trees died, trucks lugged heavy pallets of books to bookstores, more CO<sub>2</sub> made your summer hotter, and you weren't all that satisfied.

It's time to give that 19<sup>th</sup> century idea of publishing a rest. A few chapters of this book are free. You can purchase the whole book as an inexpensive weightless download, with an optional five-year subscription for updates. Why isn't it on a Kindle or an iPad? We want *you*, not Amazon or Apple, to own your copy, and we want you to update it as often as you please. You can still order a print copy. Our printing is on-demand, so your physical book was printed a few days ago instead of a few years ago, and it comes with the electronic edition and five years of updates.

From the invention of writing, through long mediaeval queues of copying monks, Gutenberg's printing press, industrial and mass-market publishing, to the Web and to current efforts to put all books ever written online, we can't be sure what will happen next—except that in a few years you'll find this book freely available on the Web just as surely as you can now find most any other book, written in English, in a library at Oxford or in the New York City Public Library. If that wrecks our current business model, so be it. Meanwhile, you are reading this because you've joined the knowledge race. There is no other marathon like it.

*Peter Brawley & Arthur Fuller*