PRENTICE HALL

UNDERSTANDING BIG DATA SCALABILITY

BIG DATA SCALABILITY SERIES, **PART I**

CORY ISAACSON

FREE SAMPLE CHAPTER

SHARE WITH OTHERS

PRAISE FOR *UNDERSTANDING BIG DATA SCALABILITY*

"This book is useful to anyone who works with data and wants to learn more about scaling. Cory helps you understand what causes databases to slow down as data volumes grow over time. He then reviews a number of strategies that you have at your disposal to manage the growth, including software and database tuning, hardware upgrades, read-replication, and ultimately horizontal partitioning of data."

—Dan Lynn, cofounder of FullContact

"*Understanding Big Data Scalability* presents the fundamentals of scaling databases from a single node to large clusters. It provides a practical explanation of what Big Data systems are, and the fundamental issues to consider when optimizing for performance and scalability. Cory draws on his many years of database experience to explain the issues involved in working with data sets that can no longer be handled with single monolithic relational databases.

"When transitioning from a traditional relational database deployment, it is tempting to ignore traditional database discipline regarding data modeling and data integrity. Much of this has been motivated by the proliferation of schema-less NoSQL databases. In spite of this trend, Cory shows why it is still important to carefully structure your data to maintain data integrity and allow sharding in such a way as to avoid costly distributed scan/shuffle operations. He discusses a practical approach to this called *relational sharding*. This is a commonsense method that avoids the pitfalls of black-box sharding. Cory's approach is particularly relevant now that relational data models are making a comeback via SQL interfaces to popular NoSQL databases and Hadoop distributions.

"*Understanding Big Data Scalability* addresses practical problems in Big Data processing systems using real-life examples. This book should be especially useful to database practitioners new to the process of scaling a database beyond a traditional single-node deployment."

—Brian O'Krafka, software architect

“Software is like magic, or so many people think. In practice, many fundamental principles of computation are constrained by immutable principles derived from the laws of physics (think transistors). One such limitation is easily understood by analyzing the trade-offs between horizontal and vertical scaling. Cory does a great job justifying the inescapable need for distributed systems, while also acknowledging their fundamental pitfalls. Good read for those who are trying to understand databases from first principles, rather than by analogy.”

—Ivan Bercovich, senior director of engineering, FindTheBest.com

“My first database project was in 1982. Since then, Moore’s law has expressed itself in just about every corner of technology—except database and database management. I’d argue alongside Cory that a significant part of the database ‘innovations’ we have experienced since then have been mostly an illusion. Can this be true? Really? If there is just a tinge of doubt in your mind, open these pages—Cory wrote this book for you.”

—Matthew Rockwell, systems design and architecture consultant

“I work with customers across the spectrum—from small start-ups to large enterprises—and they all have one thing in common: concerns and questions about how to handle database growth as their user base expands. When they have complex database scaling issues, I send them to Cory Isaacson and this book illustrates why. Not only can Cory explain the ‘whys’ behind their issues in clear language with relevant examples and analogies, but he also has the in-depth knowledge to handle the ‘hows’ in addressing those problems.”

—Brian Adler, principal cloud architect, RightScale, Inc.

UNDERSTANDING BIG DATA SCALABILITY

UNDERSTANDING BIG DATA SCALABILITY

BIG DATA SCALABILITY SERIES

PART I

Cory Isaacson

Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

ISBN-13: 978-0-13-359870-4
ISBN-10: 0-13-359870-5
Second digital release, August 2014

Executive Editor
Bernard Goodwin

Development Editor
Chris Zahn

Managing Editor
John Fuller

Full-Service Production Manager
Julie B. Nahil

Project Editor
Anna Popick

Copy Editor
Anna Popick

Indexer
Jack Lewis

Proofreader
Anna Popick

Editorial Assistant
Michelle Housley

Compositor
Anna Popick

CONTENTS

PREFACE

The *Big Data Scalability* series is a comprehensive, four-part volume containing information on many facets of database performance and scalability. This book, *Understanding Big Data Scalability*, is the first book in the series.

There is no doubt that Big Data and scalability are some of the hottest and most important topics in today's fast-growing applications. Not only is data constantly growing, the *rate* at which we accumulate data from a variety of sources is accelerating at a dramatic pace.

Understanding Big Data Scalability provides the foundation and important basic concepts for learning how scaling a Big Data infrastructure operates, and covers the most important governing factors for implementing a scalable and dynamic Big Data cluster. Later books in the series will delve into more technical detail, including the various data architectures available, optimization techniques, innovative ways to model and process data, and the types of database management system (DBMS) engines that can be used.

This book (and series) is targeted at application architects, database architects, and those with a need to scale their database tier. It is oriented toward high-volume applications.

While the subject of Big Data analytics is covered to some degree in the context of how it fits with typical online and business applications, this is not a book that concentrates exclusively on Hadoop or MapReduce batch-processing analytics. There are many books on this subject, and rather than writing another possibly redundant text, I decided to focus on the common applications we deal with every day. No doubt, these applications will require analytics capabilities, including MapReduce batch-processing, so where I discuss this approach I do so in the context of how it fits into a typical application infrastructure. In my experience, our "common" applications today are experiencing unprecedented data growth, often from multiple sources. They are driving a need for real-time processing at rates that would have been inconceivable just a short time ago. Therefore, I hope to provide you with a different perspective—one that has not been thoroughly covered by other authors.

Understanding Big Data Scalability includes the following chapters:

- Chapter 1, "Introduction": This initial chapter goes into far more detail on the purpose of the book, the origins of Big Data, and what is to be covered.
- Chapter 2, "Why Databases Slow Down": Simply put, today's typical monolithic databases frequently experience performance issues. By exploring the causes of this performance degradation, you will understand the basic principles involved that drive the need for Big Data scalability—as well as the foundation for knowing how to properly implement a scalable data architecture.
- Chapter 3, "What Is Big Data?": It is important to establish a good working definition for Big Data and how it affects your own applications.
- Chapter 4, "Big Data in the Real World": In this chapter I cover two actual case studies of common applications that had a need to scale. You will see how these solutions were implemented—information that you can use to formulate your own Big Data scalability strategy.
- Chapter 5, "Scaling Your Application": Any scalability strategy includes scaling all of your application components. That is covered here, including the various tiers of a typical application and what is involved in scaling each.
- Chapter 6, "When to Scale Your Database": Scaling your data tier is an important decision with large consequences in any application environment. It is important to know when you need to scale your database, and what factors drive scalability; these items are covered in this chapter.
- Chapter 7, "All Data Is Relational": A fundamental concept that affects all application data is data relationships. This is true regardless of the type of DBMS engine you are using, and in fact it can be even more important with non-relational databases. Thus I provide a brief overview of the importance of data relationships and how they pervade virtually everything to do with working with data.
- Chapter 8, "It's All About Sharding": In this chapter you will learn why virtually every database scalability implementation relies on some form of database sharding. You will also learn the various approaches to sharding, and the strengths and weaknesses of each.
- Chapter 9, "Scaling Big Data: The Endgame": In this final chapter of Part I of the series, I cover the most important objectives of a scalable Big Data implementation. These objectives will serve as the most important guidelines for ensuring your architecture and implementation are both successful and high performance.

There is a companion Web site for the entire series, with supplementary blog posts, forum discussions, and other relevant information related to the Big Data Scalability Series: www.bigdatascalability.com.

I hope you find this book informative and, most of all, useful, generating practical results in your own Big Data implementations. I am always interested in hearing feedback from readers, as well; you can reach me through the Big Data Scalability Series Web site noted above.

About the Author

Cory Isaacson is CEO/CTO of CodeFutures Corporation, maker of dbShards, a leading database scalability suite providing a true shared-nothing architecture for relational databases. Cory has authored numerous articles in a variety of publications, including *SOA Magazine* and *Database Trends and Applications*. He is a regular columnist for *JAXenter* magazine, and the author of the book *Software Pipelines and SOA* (Addison-Wesley, 2009). Cory has more than twenty-five years' experience with advanced software architectures, and has worked with many of the world's brightest innovators in the field of high-performance computing. Cory has spoken at hundreds of public events and seminars, and has helped numerous organizations address the real-world challenges of database performance and scalability—effectively applying Big Data technologies to a variety of fields, including social networking, mobile applications, gaming, and high-volume transaction systems. In his prior position as president of Rogue Wave Software, he actively led the company back to a position of profitable growth, culminating in a successful acquisition by a leading private equity firm. Cory can be reached through the Big Data Scalability Series Web site: www.bigdatascalability.com.

1 INTRODUCTION

Big Data Scalability is a comprehensive, four-part volume containing information on many facets of database performance and scalability. I have found through years of experience that in order to master any aspect of software, databases, or computer science, the most important thing is a solid foundation and understanding of the fundamental concepts. Once you have this knowledge, you can apply it to solve real-world problems with your own database tier.

This series is primarily focused on how to apply Big Data scalability to common application problems with the wide selection of DBMS engines available today. As such, it is less focused on pure analytics Big Data problems with engines like Apache Hadoop (there are many other books dedicated to that), and more directed toward taking advantage of Big Data scalability with popular DBMS engines used in enterprise and online applications.

WHAT YOU WILL LEARN

This part of the series will cover the following critical topics for building a high-performance, scalable, and reliable database tier:

- The primary causes of database performance issues
- What Big Data is

- Options for scaling your application
- When and how to scale your database
- The fundamentals of data distribution (database sharding)

THE CHALLENGE OF BIG DATA

Have you ever experienced a database performance problem? It's a silly question, I know. Everyone who has worked with databases for any amount of time has seen performance issues. In fact, it's so common that *database performance* is almost an oxymoron. The very nature of database technology is to get slower over time, sometimes dramatically so. Why? Because databases only tend to get bigger through increased usage, and of course with more data the database management system (DBMS) has to do more work to keep up with application demands. Add to the mix an explosive high-transaction volume and you have the perfect recipe for something we all dread: *slow database performance*. In fact, sometimes it's downright *horrible* database performance.

What about the memorable experience of a "database down" on a live production system? This can easily be your worst day on the job. Sometimes, when you get that call, you wish you could just forget it and not even go to work. Not only are big, growing databases hard to keep fast, they are very painful to repair when something bad happens. I've seen applications that were down for days due to extended database recovery times, and I'm sure you have heard of some notorious all-too-public examples of this as well. So an all-important part of a Big Data success is, in fact, planning for failure—ensuring you are prepared, with a robust, high-availability (HA) runtime environment and a solid disaster recovery (DR) plan for a worst-case scenario (they do happen).

What can you do about it? That's what this book is all about: what *you* can do to make your database tier fast, scalable, and reliable, delivering the performance that your application needs. The focus of this portion of the book is not on any particular DBMS or database technology, but rather the answers to these questions—questions that apply to any database and any DBMS technology:

- What are the primary causes of database slow-down?
- How do you keep your database tier reliable and operational (especially in demanding, 24X7 environments)?
- Are there short-term, quick fixes you can implement to ease the pain?

- What are the best ways to scale your database to accommodate extreme growth in transaction volume and data sizes?
- How do you pick the best database technology (often *technologies*) for your application needs?

The answers to these questions and more are covered throughout the text. The aim is to arm you with the important concepts and facts common to all DBMS types, so that you can make the best decisions possible in managing your own data explosion, and come out on top.

Please note that in subsequent parts of the series I will provide a tour and overview of many common and leading-edge DBMS engines, both conceptually by category as well as specific solutions you can use in your Big Data arsenal. With this information, combined with the fundamentals covered in Part I, my hope is to assist you in preparing for your own Big Data adventure—commensurate with much success and reward.

TODAY'S BIG DATA EXPLOSION

Next let's look at the drivers and scope of Big Data.

MANAGING AND CAPITALIZING ON THE CURRENT DATA BOOM

We are living in the midst of a *data explosion*, a true boom in databases and database technology, the likes of which the world has never seen. Since you are reading this book, I assume you have an interest in database performance and scalability, the same as I do. Figure 1.1 illustrates the Big Data explosion by the current data boom, and how critical it is for us to be able to extract *meaning* from all of this data.

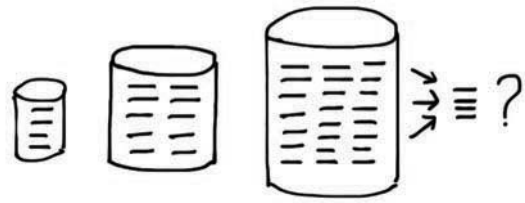

Figure 1.1 The Big Data explosion

What do I mean by a *data boom*? Given that information is often the most valuable commodity in today's tech-centric world, this means that we as data professionals

hold the keys to the kingdom. Not only is data getting more massive each and every year, the rate at which data is being generated is accelerating at a super-linear rate. It's tough enough to figure out how to capture and manage all that data, but more importantly we need to identify the right set of architectural decisions, tools, and capabilities to allow our organizations to capitalize on that data. In other words, giving meaning to raw data is just as important as the collection, reliability, and management of Big Data environments.

The causes of the explosion of database data are not hard to find:

- The advent of the Internet and the World Wide Web has generated exponential growth in the global user community—users with ever-expanding access to computing power and bandwidth.
- The interaction of these users with Internet applications has resulted in unprecedented levels of data and transaction volumes.
- The shift to online advertising supported by the likes of Google, Yahoo, and others is a key driver in the data boom we are seeing today.
- The overall expansion of the worldwide economy has spurred massive data growth for traditional commerce (e.g., increased airline travel, international purchases, online products, etc.).
- The core social networks (e.g., Facebook, Twitter, LinkedIn, and now Google+), by their very nature, have generated massive new ways for people to communicate and interact, resulting in correspondingly large data sets and transaction volume.
- Many specialized social networks have also arisen—everything from match-making sites to special interest groups, and even "buy-sell" applications that have generated their own micro-economies.
- An entirely new breed of social network applications has been spawned, leveraging the inter-connection of social network users in fascinating ways, driving exponential growth in application volume, again with huge transaction volumes and data sizes (sometimes virtually overnight success stories).
- Web- and advertising-analytics applications abound, crawling and analyzing virtually every aspect of the user interaction described above, again resulting in massive data sets with intense database access needs.
- An entirely new breed of chatter trend analytics applications have emerged, analyzing things like Twitter tweets, Facebook chats, and so on, requiring massive levels of data storage and access.

- Last, the world has gone mobile. In fact, in burgeoning economies and established countries alike, smart phones and tablets are by far the most readily available, high-growth, and commonly used communication vehicle for much of the world's population, generating a nearly incomprensible stream of data, transactions, application interaction, and messaging volume (with no end in sight).

Any one of the factors above would, in itself, have created unprecedented growth in databases, but taken collectively the impact is truly awe-inspiring and overwhelming. I never cease to be amazed by the number of extremely bright people I meet that have an ever-widening set of ideas—ideas that are building new businesses that are further fueling the growth of databases to support these fast-growing industries and sectors.

Your Role as a Data Architect

For those of us who are data architects, this is the most exciting time ever. Make no mistake about it: *You* are the key to the success of this new environment. For without the talent and dedication of the individual data architect to find solutions to address today's data explosion, none of these ventures would succeed.

If you are a database architect or database administrator (DBA)—one who is driven to deliver the fastest and most scalable database platform possible for your organization—this book is addressed to you. In the end, it is your ability and intelligence that will solve these problems for real; all the technology in the world can at best assist you to do the job.

The Acceleration of Big Data Innovation

The phenomenon of the Big Data explosion has not only yielded huge growth in data, but also has spurred ingenious innovation to conquer the Big Data challenge.

The sheer number of new database platforms and DBMS engines introduced in the past few years is mind-boggling. Daunting as it may seem to stay current with this rate of technological advancement, the innovation itself is nothing but good for the professional data architect. This investment in Big Data solutions has yielded an incredibly wide array of technology options for you to consider.

While the sheer rate of innovation has created an environment that is often nothing short of exhilarating (at least for a database fanatic like me), it can at times be intimidating as well.

Therefore, knowing how to evaluate these technologies so you can pick the right options for your application is critical. In fact, it's a vital skill you (and the companies or organizations you support) cannot afford to be without. A portion of this series is devoted to a review of the various types of data platforms available, explaining typical use cases, available solutions, and how and where to apply them. But more importantly, the underlying concepts—applicable to all DBMS engines—are covered, enabling you to make informed evaluations and decisions for your specific application requirements.

I firmly believe (and have found through hard-won experience) that the more you understand the fundamentals of Big Data platforms and technologies, the better you will be at implementing truly successful database environments, taming whatever Big Data explosion you are in the midst of.

BACKGROUND FOR THIS BOOK

My favorite subject is databases—pretty much everything about them, from database performance to management and of course scalability most of all. You might say I'm a bit of fanatic about it (okay, maybe I'm an extreme fanatic), but I wouldn't do it if I didn't love the field so much. I have been working intensively with all types of database technologies throughout my career. In one of my first companies, we built a successful independent database and application development firm, beginning with Sybase and expanding into other DBMS engines after that. More recently I ran Rogue Wave Software, a C++ tools vendor with a large concentration of Wall Street customers. While working with these customers, which included some of the top architects in the world, I saw a clear trend that spurred my interest even further: No matter how well an application could be scaled, the database tier was *always* the bottleneck. In fact, I learned that the database tier forms the *last mile* of application scalability, being at the very foundation of all functionality. The message was clear: Find a solution to that problem, and the potential was huge.

I have experienced the very best (and the very, very worst) that database management systems have to offer, in the trenches and firsthand. I've had the fortune and privilege of serving as the architect and manager for many mission-critical applications, from the advent of client-server technology to the Big Data systems of today. In addition to having lived through just about every database nightmare you can imagine, I've also seen the results when the database tier is performing outrageously well. I've seen how hard great software teams work to conquer these problems; there are few careers that require as much intelligence,

persistence, and long, demanding hours. This has given me a profound respect for sharp technologists that conquer Big Data problems, the engines that are the very backbone of every application, not to mention the foundation of much of today's economic growth. In fact, it's obvious that the database explosion of the information age is upon us, and it's likely to continue for many years to come as new and clever business opportunities arise to capitalize on this first-ever phenomenon encountered in world history.

Through this extensive and often intense experience, I've gained many insights over the years. My intent in writing this book is to share this experience with a wide audience, with the hope that you can avoid many of the common pitfalls I have seen repeated over and over. The book is written from a practical point of view, with what I believe are the salient points required to make the database tier of an application really cook. Many of the techniques can be applied easily, without the need to purchase or implement new technology, while others are intended as guides for taking full advantage of the latest and greatest database capabilities emerging from the fantastic technical advances made available in the past few years.

As CEO and CTO of CodeFutures (the makers of dbShards and AgilData technology), I have had the pleasure of working with the sharpest technical teams in my career—not just within our own engineering team, but also the customers that demanded the absolute best in database performance and reliability. We have worked with some of the world's largest social networking applications, with phenomenal data growth and transaction volumes. Frankly, in the early days of our technology, we lived through a few genuinely hair-raising experiences, but with hard work in partnership with our customers, we always were able to succeed, even against the toughest challenges. More important, the wide array of applications where we have applied database sharding resulted in a broad range of experience, with large data sets and transaction volumes that were unimaginable just a few years ago. Many of these applications experienced extreme growth in transaction volume, an extraordinary thing to behold (to put it mildly). For example, watching a social game grow from nothing to over 1 million daily users in a few short months is a sight to see, and when you are part of the team responsible for ensuring the database platform handles the load, continuously running 24X7, it's an experience that really grabs your attention (and can result in some very late nights!).

As a disclaimer right up front, many of the examples and experiences are derived from our work with our own products. dbShards, and our newest offering, AgilData, are most definitely commercial products, and of course part of my purpose in writing this book is to gain more attention for what we deliver. Having said that, I have worked hard to keep the book balanced, concentrating on the important facts

that drive database scalability irrespective of a particular DBMS engine or technology. In fact, the book discloses a lot of the "tricks" and designs we use to deliver such scalable performance, with the hope that you will find them beneficial and useful in your own database environment. And don't worry, we have new capabilities coming out with our advanced AgilData technology that will blow the world away. Who knows, they may provide material for future books in this series.

WHY THE FOCUS ON DATABASE SHARDING?

You will notice as you read this book that much of the focus is on database sharding, a technique and architecture for horizontal *shared-nothing* partitioning of database data across independent nodes or servers. From extensive experience and research, we have found that, regardless of what DBMS you use, database sharding is the only effective means for scaling a database. This applies to all types of databases, including the traditional relational database management system (RDBMS), NoSQL platforms, Big Data engines, Cloud databases, or databases of the so-called NewSQL paradigm. If you look closely at any of these offerings, particularly if they advertise a scalable platform, you will see that some sort of horizontal partitioning of data across nodes is used. Thus, database sharding is the answer, used in one way or another for virtually every leading database platform. Figure 1.2 shows the concept of database sharding, breaking a large monolithic database into multiple, smaller, sharded database instances across multiple servers.

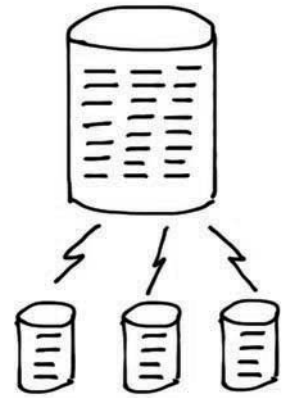

Figure 1.2 Database sharding

Understanding the underlying principles of how database sharding works is the answer to controlling your own destiny regarding the ultimate performance of your

Big Data implementation. The basic concepts behind sharding and how it works are not difficult, and once you grasp them you can implement amazing levels of scalability. Ultimately this book is about making you successful with your database tier, and the more you know about this concept, the more power you will have over defining, predicting, optimizing, and controlling the performance of your own databases.

NOTE

Some vendors have downplayed the importance of database sharding, or have outright stated that sharding is an outmoded or useless idea. Nothing could be further from the truth, and once you understand the full meaning and applicability of the term, you will see that these very same vendors are using database sharding in their own technologies. Much more on the topic will be covered throughout the series.

SUMMARY

I hope you find the ensuing pages interesting, informative, and most of all useful. If you gain even a single new idea that results in improved database performance for your application, I will have achieved what I set out to do with this work.

INDEX

M

N

O

P

Q

R

S

T

U

V

W