



HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE



**DOWNLOAD EBOOK : HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE
PDF**



Random Access to Your Planet-Size Data



HBase

The Definitive Guide

O'REILLY®

Lars George

Click link below and free register to download ebook:
HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE

[DOWNLOAD FROM OUR ONLINE LIBRARY](#)

HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE PDF

HBase: The Definitive Guide By Lars George Exactly how can you transform your mind to be more open? There lots of sources that could help you to enhance your thoughts. It can be from the other experiences and tale from some people. Reserve HBase: The Definitive Guide By Lars George is one of the relied on resources to obtain. You can discover plenty books that we discuss right here in this web site. As well as now, we reveal you among the most effective, the HBase: The Definitive Guide By Lars George

About the Author

Lars George has been involved with HBase since 2007, and became a full HBase committer in 2009. He has spoken at various Hadoop User Group meetings, as well as large conferences such as FOSDEM in Brussels. He also started the Munich OpenHUG meetings. He now works closely with Cloudera to support Hadoop and HBase in and around Europe through technical support, consulting work, and training.

HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE PDF

[Download: HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE PDF](#)

Some people could be chuckling when looking at you checking out **HBase: The Definitive Guide By Lars George** in your downtime. Some could be admired of you. And also some may desire resemble you that have reading leisure activity. What about your personal feel? Have you felt right? Reading HBase: The Definitive Guide By Lars George is a requirement and also a leisure activity at the same time. This problem is the on that will certainly make you really feel that you should check out. If you know are trying to find the book entitled HBase: The Definitive Guide By Lars George as the selection of reading, you could discover right here.

It is not secret when attaching the writing abilities to reading. Checking out *HBase: The Definitive Guide By Lars George* will certainly make you get even more sources and also resources. It is a way that could boost just how you neglect and also understand the life. By reading this HBase: The Definitive Guide By Lars George, you could greater than exactly what you obtain from various other publication HBase: The Definitive Guide By Lars George This is a well-known book that is published from renowned author. Seen kind the author, it can be relied on that this publication HBase: The Definitive Guide By Lars George will offer lots of motivations, concerning the life as well as experience and everything within.

You may not have to be doubt concerning this HBase: The Definitive Guide By Lars George It is easy way to obtain this book HBase: The Definitive Guide By Lars George You could just see the distinguished with the web link that we offer. Below, you could acquire the book HBase: The Definitive Guide By Lars George by on-line. By downloading HBase: The Definitive Guide By Lars George, you can discover the soft data of this book. This is the exact time for you to start reading. Also this is not printed book HBase: The Definitive Guide By Lars George; it will specifically offer even more perks. Why? You could not bring the printed book HBase: The Definitive Guide By Lars George or only pile the book in your property or the workplace.

HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE PDF

If you're looking for a scalable storage solution to accommodate a virtually endless amount of data, this book shows you how Apache HBase can fulfill your needs. As the open source implementation of Google's BigTable architecture, HBase scales to billions of rows and millions of columns, while ensuring that write and read performance remain constant. Many IT executives are asking pointed questions about HBase. This book provides meaningful answers, whether you're evaluating this non-relational database or planning to put it into practice right away.

- Discover how tight integration with Hadoop makes scalability with HBase easier
 - Distribute large datasets across an inexpensive cluster of commodity servers
 - Access HBase with native Java clients, or with gateway servers providing REST, Avro, or Thrift APIs
 - Get details on HBase's architecture, including the storage format, write-ahead log, background processes, and more
 - Integrate HBase with Hadoop's MapReduce framework for massively parallelized data processing jobs
 - Learn how to tune clusters, design schemas, copy tables, import bulk data, decommission nodes, and many other tasks
-
- Sales Rank: #610897 in Books
 - Published on: 2011-09-23
 - Original language: English
 - Number of items: 1
 - Dimensions: 9.19" h x 1.17" w x 7.00" l, 1.88 pounds
 - Binding: Paperback
 - 556 pages

About the Author

Lars George has been involved with HBase since 2007, and became a full HBase committer in 2009. He has spoken at various Hadoop User Group meetings, as well as large conferences such as FOSDEM in Brussels. He also started the Munich OpenHUG meetings. He now works closely with Cloudera to support Hadoop and HBase in and around Europe through technical support, consulting work, and training.

Most helpful customer reviews

25 of 26 people found the following review helpful.

Probably a great book a couple of years ago

By yepher

I am just getting started with HBase and was looking for a book that gets me up and running. I read through the first half of the book and was getting a pretty good feel for how it seems you should use HBase.

Unfortunately the API used are out of date. I was unable to get any of the examples to work with HBase

0.94.4. I tried downgrading HBase to 0.92.0 but it did not work either. The Maven pom files make it look easy to up the version but it turns out the API has changed quite a bit.

It would be really nice if the GitHub project was brought up to date so one reading through the book can at least get the examples to run with a current version of HBase. I tried several of the project forks but none seemed to work for me. The project README should at least be updated to point to the specific version of HBase that is needed to run the examples.

I made it to chapter four when I realized this book is mostly useless at this point.

15 of 19 people found the following review helpful.

Not Likely to be Surpassed

By Glenn Street

When a book bills itself as "The Definitive Guide," well, that's a tall order to fill. But, except for updates as new releases of HBase roll out, I can't imagine another book surpassing this one by Lars George.

Lars George has been working with HBase since 2007 and is a full committer to the project as of 2009. He now works for Cloudera (a company providing a commercial flavor of Hadoop, as well as Hadoop support). After reading this book, there's no question in my mind that George has deep understanding, not only of HBase as a data solution, but of the internal workings of HBase.

George gives the background and history of HBase in the larger context of relational databases and NoSQL, which I found to be very helpful. The many diagrams throughout the book are extremely useful in explaining concepts, especially for those of us coming from a relational database background.

George has an excellent and clear writing style. Take, for example, the section where he discusses The Problem with Relational Database Systems, giving a quick rundown of the typical steps for getting an RDBMS to scale up. The flow of his summary reads like the increasing levels of panic that many of us have gone through when dealing with a database-backed application that will not scale.

As an example of how thorough and comprehensive the book is, look at chapter 2, where there is an extensive discussion of the type and class (not desktop PCs!) of machines suitable for running HBase. George gives a truly helpful set of configuration practices, even down to a recommendation for having redundant power supply units.

Another example of his thoroughness comes where George discusses delete methods (Chapter 3). He shows how you can use custom versioning, while admitting that the example is somewhat contrived. Indeed, right after elaborating the example, there is a distinct "Warning" box that admits that custom versioning is not actually recommended. So, even though you may not implement custom versioning, you do understand it as a feature that HBase provides.

Many of the programming examples come with excellent remarks or discussions of the tradeoffs implicit in the techniques, including performance and scaling concerns. Java developers will be most comfortable with the majority of examples, but they can be followed by anyone with some object-oriented programming experience.

I really appreciated the thorough discussion in chapter 8 ("Architecture") of subjects like B+ trees vs. Log-Structured Merge Trees (LSMs), the Write-Ahead Log, and seeks vs. transfers, topics which are relevant not only to HBase but to many database systems of varying architectures.

The level of thoroughness is also the book's only weakness. I'm not sure who the target audience for this book is, because it serves both developers and system or database administrators. While nearly every imaginable HBase topic is touched upon, some would have been better off merely listed, with appropriate references given to sources of more information (for example, all those hardware recommendations). The print edition of the book is 552 pages.

Still, a complaint that a book is too detailed shouldn't be interpreted as much of a complaint. Anyone with an interest in NoSQL databases in general, and HBase in particular should read and study this book. It's not likely to be superseded in the future.

5 of 7 people found the following review helpful.

Wide scope and deep coverage by the "Ask Tom" of HBase

By Erik Gfesser

This text is amongst the few books I have read in my career which not only serves as a great introduction to a technology, but also provides significant content for both advanced users and those who wish to understand what is under the covers in order to understand both benefits and drawbacks before making a decision in terms of adoption. While I was not familiar with the author, Lars George, until reading this text, the foreword to what the author shares here offers some street cred. "Of all those who have contributed to HBase over the years, it is poetic justice that Lars is the one to write this book. Lars was always dogging HBase contributors that the documentation needed to be better if we hoped to gain broader adoption. Everyone agreed, nodded their heads in ascent, amen'd, and went back to coding. So Lars started writing critical how-tos and architectural descriptions inbetween jobs and his intra-European travels as unofficial HBase European ambassador. His Lineland blogs on HBase gave the best description, outside of the source, of how HBase worked, and at a few critical junctures, carried the community across awkward transitions (e.g. an important blog explained the labyrinthian HBase build during the brief period we thought an Ivy-based build to be a 'good idea'). His luscious diagrams were poached by one and all wherever an HBase presentation was given."

At least one other reviewer here commented that it was hard for them to understand the purpose of this book, because of its broad scope. But in my opinion its broad scope is what makes it so valuable, at least to the architect. The potential reader just needs to keep in mind that they need to understand what they wish to gain from reading a text on HBase. For example, if the potential reader simply wants to understand how to use the APIs that HBase provides, there may be no need to understand HBase architecture and so the chapters that the author provides on related material may not be applicable at the outset. It is difficult to argue for a sequential reading of this text from beginning to end, although the potential reader also needs to keep in mind the importance of minimal understanding of the programmatic connections with the architecture, in which case a cursory overview of some of the more advanced material is highly recommended. In contrast, as an architect who wants to understand as many of the benefits, drawbacks, and implementation details as possible, especially when considering my extensive experience with relational databases such as Oracle, I appreciate all of the detail that George has to offer here, and would not have it any other way. While I have not yet reviewed his blog yet to any great extent, based on his street cred and the abundance of material in this book, I do not think it is a stretch to akin him to Tom Kyte of "Ask Tom" fame within the Oracle space.

After an introduction that provides discussions on big data, column-oriented databases, problems with relational database systems, nonrelational database systems, and an HBase architectural overview (all within Chapter 1), George quickly moves forward to a chapter on HBase installation (Chapter 2), followed by discussions of native Java APIs (Chapters 3, 4, and 5), available clients (Chapter 6), and integration with Hadoop's MapReduce framework (Chapter 7). The remainder of the text delves into more advanced topics, beginning with HBase architecture (Chapter 8), followed by design topics critical to taking full advantage of HBase architecture (Chapter 9), as well as discussions on cluster monitoring (Chapter 10), performance

tuning (chapter 11), and cluster administration (chapter 12). As a general rule, George does a great job at introducing topics in an order that makes sense without frequently referring to chapters later in the book, which is a pet peeve of mine with other technical books. One of the many aspects I especially appreciated about this text is that even in chapters that cover content that might be seen as high-level content by other architects, George is not afraid to get down in the trenches at the programmatic level, even in the preface to the book, in which the author explains how to build the examples that he provides in his Git repository.

Please be aware that even though the author discusses features such as coprocessors that have been included in subsequent releases (at the time of writing this review, version 0.92.1 has been made available as a stable release), the author does make note in Chapter 2 that "the current version of HBase will only run on Hadoop 0.20.x. It will not run on Hadoop 0.21.x (nor 0.22.x) as of this writing. HBase may lose data in a catastrophic event unless it is running on an HDFS that has durable sync support. Hadoop 0.20.2 and Hadoop 0.20.203.0 do not have this support. Currently, only the branch-0.20-append branch has this attribute. No official releases have been made from this branch up to now, so you will have to build your own Hadoop from the tip of this branch." The author does bring up the fact that Cloudera's CDH3 has already applied the patches associated with durable sync, if building Hadoop is not your cup of tea, but keep in mind that Cloudera as of less than two weeks ago has already released CDH4 Beta 2. Anyone interested in HBase needs to keep in mind that this space is moving along quickly, so plan to get started with this book and move forward with the abundant material that is increasingly being made available online. That said, George provides a wealth of knowledge that I do not expect to be surpassed at any time in the foreseeable future by any other publication.

See all 14 customer reviews...

HBASE: THE DEFINITIVE GUIDE BY LARS GEORGE PDF

You can finely include the soft file **HBase: The Definitive Guide By Lars George** to the gadget or every computer hardware in your office or residence. It will assist you to always continue reviewing HBase: The Definitive Guide By Lars George each time you have extra time. This is why, reading this HBase: The Definitive Guide By Lars George doesn't offer you troubles. It will give you important resources for you which wish to start writing, discussing the comparable book HBase: The Definitive Guide By Lars George are various publication field.

About the Author

Lars George has been involved with HBase since 2007, and became a full HBase committer in 2009. He has spoken at various Hadoop User Group meetings, as well as large conferences such as FOSDEM in Brussels. He also started the Munich OpenHUG meetings. He now works closely with Cloudera to support Hadoop and HBase in and around Europe through technical support, consulting work, and training.

HBase: The Definitive Guide By Lars George Exactly how can you transform your mind to be more open? There lots of sources that could help you to enhance your thoughts. It can be from the other experiences and tale from some people. Reserve HBase: The Definitive Guide By Lars George is one of the relied on resources to obtain. You can discover plenty books that we discuss right here in this web site. As well as now, we reveal you among the most effective, the HBase: The Definitive Guide By Lars George