1.	Record Nr.	UNINA9910792825603321
	Autore	Yarabarla Sandeep
	Titolo	Learning apache cassandra : managing fault-tolerant and scalable data / / Sandeep Yarabarla
	Pubbl/distr/stampa	Birmingham, England ; ; Mumbai, [India] : , : Packt, , 2017 ©2017
	ISBN	1-78712-840-7
	Edizione	[Second edition.]
	Descrizione fisica	1 online resource (354 pages) : illustrations
	Disciplina	005.3
	Soggetti	Database management Database design Web services Information storage and retrieval systems
	Lingua di pubblicazione	Inglese
	Formato	Materiale a stampa
	Livello bibliografico	Monografia
	Note generali	Includes index.
	Sommario/riassunto	Build a scalable, fault-tolerant and highly available data layer for your applications using Apache Cassandra About This Book Install Cassandra and set up multi-node clusters Design rich schemas that capture the relationships between different data types Master the advanced features available in Cassandra 3.x through a step-by-step tutorial and build a scalable, high performance database layer Who This Book Is For If you are a NoSQL developer and new to Apache Cassandra who wants to learn its common as well as not-so-common features, this book is for you. Alternatively, a developer wanting to enter the world of NoSQL will find this book useful. It does not assume any prior experience in coding or any framework. What You Will Learn Install Cassandra Create keyspaces and tables with multiple clustering columns to organize related data Use secondary indexes and materialized views to avoid denormalization of data Effortlessly handle concurrent updates with collection columns Ensure data integrity with lightweight transactions and logged batches Understand eventual consistency and use the right consistency level for your situation Understand data distribution with Cassandra Develop simple application using Java driver and implement

application-level optimizations In Detail Cassandra is a distributed database that stands out thanks to its robust feature set and intuitive interface, while providing high availability and scalability of a distributed data store. This book will introduce you to the rich feature set offered by Cassandra, and empower you to create and manage a highly scalable, performant and fault-tolerant database layer. The book starts by explaining the new features implemented in Cassandra 3.x and get you set up with Cassandra. Then you'll walk through data modeling in Cassandra and the rich feature set available to design a flexible schema. Next you'll learn to create tables with composite partition keys, collections and user-defined types and get to know different methods to avoid denormalization of data. You will then proceed to create user-defined functions and aggregates in Cassandra. Then, you will set up a multi node cluster and see how the dynamics of Cassandra change with it. Finally, you will implement some application-level optimizations using a Java client. By the end of this book, you'll be fully equipped to build powerful, scalable Cassandra database layers for your applications. Style and approach This book takes a step-...