

1. Record Nr.	UNINA9910808956903321
Autore	Mohan Surendra
Titolo	Apache Solr high performance : boost the performance of Solr instances and troubleshoot real-time problems // Surendra Mohan ; cover image by Glain Clarrie
Pubbl/distr/stampa	Birmingham, England : , : Packt Publishing Ltd, , 2014 ©2014
ISBN	1-78216-483-9
Descrizione fisica	1 online resource (124 p.)
Collana	Community Experience Distilled
Disciplina	005.758
Soggetti	Search engines - Programming Open source software
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Cover; Copyright; Credits; About the Author; About the Reviewers; www.PacktPub.com; Table of Contents; Preface; Chapter 1: Installing Solr; Prerequisites; Installing components; Summary; Chapter 2: Boost Your Search; Scoring; Boosting query-time and index-time; Index-time boosting; Query-time boosting; Troubleshoot queries and scores; The dismax query parser; Lucene DisjunctionMaxQuery; Autophrase boosting; Configuring autophrase boosting; Configuring the phrase slop; Boosting a partial phrase; Boost queries; Boost functions; Boost addition and multiplication; Function queries Field referencesFunction references; Mathematical operations; The ord() and rord() functions; Other functions; Boosting the function query; Logarithm; Reciprocal; Linear; Inverse reciprocal; Summary; Chapter 3: Performance Optimization; Solr performance factors; Solr caching; Document caching; Query result caching; Filter caching; Result pages caching; Using SolrCloud; Creating a SolrCloud cluster; Multiple collections within a cluster; Managing a SolrCloud cluster; Distributed indexing and searching; Stopping automatic document distribution; Near real-time search; Summary Chapter 4: Additional Performance Optimization TechniquesDocuments similar to those returned in the search result; Sorting results by function values; Searching for homophones; Ignore the defined words

from being searched; Summary; Chapter 5: Troubleshooting; Dealing with the corrupt index; Reducing the file count in the index; Dealing with the locked index; Truncating the index size; Dealing with a huge count of open files; Dealing with out-of-memory issues; Dealing with an infinite loop exception in shards; Dealing with expensive garbage collection

Bulk updating a single field without full indexationSummary; Chapter 6: Performance Optimization with ZooKeeper; Getting familiar with ZooKeeper; Prerequisites for a distributed server; Aid your distributed system using ZooKeeper; Setting an ideal node count for ZooKeeper; Setting up, configuring, and deploying ZooKeeper; Setting up ZooKeeper; Configuring Zookeeper; Deploying ZooKeeper; Applications of ZooKeeper; Summary; Appendix; Index

Sommario/riassunto

This book is an easy-to-follow guide, full of hands-on, real-world examples. Each topic is explained and demonstrated in a specific and user-friendly flow, from search optimization using Solr to Deployment of Zookeeper applications. This book is ideal for Apache Solr developers and want to learn different techniques to optimize Solr performance with utmost efficiency, along with effectively troubleshooting the problems that usually occur while trying to boost performance. Familiarity with search servers and database querying is expected.
