

1. Record Nr.	UNINA9910483231603321
Autore	Harrison Guy
Titolo	MongoDB performance tuning : optimizing MongoDB databases and their applications / / Guy Harrison, Michael Harrison
Pubbl/distr/stampa	[Place of publication not identified] : , : APress, , [2021] ©2021
ISBN	1-4842-6879-2
Descrizione fisica	1 online resource (xxiv, 349 pages) : illustrations
Disciplina	005.7565
Soggetti	Database management
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Intro -- Table of Contents -- About the Authors -- About the Technical Reviewer -- Acknowledgments -- Introduction -- Part I: Methods and Tools -- Chapter 1: Methodical Performance Tuning -- A Cautionary Tale -- Symptomatic Performance Tuning -- Systematic Performance Tuning -- Anatomy of a Database Request -- The Layers of a MongoDB Database -- Minimizing the Application Workload -- Reducing Physical IO -- Optimizing Disk IO -- Cluster Tuning -- Summary -- Chapter 2: MongoDB Architecture and Concepts -- The MongoDB Document Model -- JSON -- Binary JSON (BSON) -- Collections -- MongoDB Schemas -- The MongoDB Protocol -- Wire Protocol -- MongoDB Drivers -- MongoDB Commands -- The find Command -- The aggregate Command -- Data Manipulation Commands -- Consistency Mechanisms -- Read Preference and Write Concern -- Transactions -- Query Optimization -- MongoDB Architecture -- Mongod -- Storage Engines -- Replica Sets -- Sharding -- Sharding Mechanisms -- Cluster Balancing -- Conclusion -- Chapter 3: Tools of the Trade -- Introduction to explain() -- Getting Started with explain() -- Alternate Plans -- Execution Statistics -- Using explain() to Tune a Query -- Visual Explain Utilities -- The Query Profiler -- The system.profile Collection -- Analyzing Profiling Data -- Tuning with MongoDB Logs -- Server Statistics -- Examining Current Operations -- Operating System Monitoring -- MongoDB Compass -- Summary -- Part II: Application and Database Design -- Chapter 4: Schema Modelling -- The Guiding

Principles -- Linking vs. Embedding -- A Case Study -- Getting All the Data for a Customer -- Fetching All Open Orders -- Top Products -- Inserting New Orders -- Updating Products -- Deleting a Customer -- Case Study Summary -- Advanced Patterns -- Subsetting -- Vertical Partitioning -- The Attribute Pattern -- Summary -- Chapter 5: Indexing -- B-Tree Indexes. Index Selectivity -- Unique Indexes -- Index Scans -- Case-Insensitive Searches -- Compound Indexes -- Compound Index Performance -- Compound Index Key Order -- Guidelines for Compound Indexes -- Covering Indexes -- Index Merges -- Partial and Sparse Indexes -- Partial Indexes -- Sparse Indexes -- Using Indexes for Sorting and Joining -- Sorting -- Using Indexes for Joins -- Index Overhead -- Wildcard Indexes -- Text Indexes -- Text Index Performance -- Geospatial Indexes -- Geospatial Index Performance -- Geospatial Index Limitations -- Summary -- Part III: Tuning MongoDB Code -- Chapter 6: Query Tuning -- Caching Results -- Optimizing Network Round Trips -- Projections -- Batch Processing -- Avoiding Excessive Network Round Trips in Code -- Bulk Inserts -- Application Architecture -- Choosing an Index vs. a Scan -- Overriding the Optimizer with Hints -- Optimizing Sort Operations -- Picking or Creating the Right Index -- Filter Strategies -- Not Equals Conditions -- Range Queries -- OR or IN Operations -- Array Queries -- Regular Expressions -- exists Queries -- Optimizing Collection Scans -- Summary -- Chapter 7: Tuning Aggregation Pipelines -- Tuning Aggregation Pipelines -- Optimizing Aggregation Ordering -- Automatic Pipeline Optimizations -- Optimizing Multi-collection Joins -- Join Order -- Optimizing Graph Lookups -- Aggregation Memory Utilization -- Sorting in Aggregation Pipelines -- Indexed Aggregation Sorts -- Disk Sorts -- Optimizing Views -- Materialized Views -- Summary -- Chapter 8: Inserts, Updates, and Deletes -- Fundamentals -- Filter Optimizations -- Explaining a Data Manipulation Statement -- Index Overhead -- Finding Unused Indexes -- Write Concern -- Inserts -- Batch Processing -- Cloning Data -- Loading from Files -- Updates -- Dynamic Value Bulk Updates -- The multi:true Flag -- Upserts -- Bulk Upsert with merge. Delete Optimizations -- Summary -- Chapter 9: Transactions -- Transaction Theory -- MongoDB Transactions -- Transaction Limits -- TransientTransactionErrors -- Transactions in the MongoDB Drivers -- The Performance Implications of TransientTransactionErrors -- Transaction Optimization -- Avoiding Transactions -- Ordering of Operations -- Partitioning Hot Documents -- Conclusion -- Chapter 10: Server Monitoring -- Host-Level Monitoring -- Network -- CPU -- Memory -- Disk IO -- MongoDB Server Monitoring -- Compass -- Free Monitoring -- Ops Manager -- MongoDB Atlas -- Third-Party Monitoring Tools -- Summary -- Part IV: Server Tuning -- Chapter 11: Memory Tuning -- MongoDB Memory Architecture -- Host Memory -- Measuring Memory -- WiredTiger Memory -- Cache Size -- Determining the Optimum Cache Size -- The Database Cache "Hit" Ratio -- Evictions -- Blocking Evictions -- Checkpoints -- WiredTiger Concurrency -- Reducing Application Memory Demand -- Document Design -- Indexing -- Transactions -- Summary -- Chapter 12: Disk IO -- IO Fundamentals -- Latency and Throughput -- Queuing -- Sequential and Random IO -- Disk Hardware -- Magnetic Disks (HDD) -- Solid State Drives -- SSD Storage Hierarchy -- Write Performance -- Write Endurance -- Garbage Collection and Wear Levelling -- SATA vs. PCI -- Recommendations for SSDs -- Storage Arrays -- RAID Levels -- The RAID 5 Write Penalty -- Non-volatile Caches in RAID 5 Devices -- Do It Yourself Arrays -- Hardware Storage Arrays -- Cloud Storage --

Disk Devices in MongoDB Atlas -- MongoDB IO -- Temporary File IO --
The Journal -- Moving the Journal to a Dedicated Device -- Datafile IO
-- Datafile Writes -- Splitting Up Datafiles Across Multiple Devices --
Detecting and Solving IO Problems -- Increasing IO Subsystem
Bandwidth -- Dedicated Server with Dedicated Disks -- Storage Arrays
-- Cloud Storage -- MongoDB Atlas.
Summary -- Chapter 13: Replica Sets and Atlas -- Replica Set
Fundamentals -- Using Read Preference -- Setting Read Preference --
maxStalenessSeconds -- Tag Sets -- Write Concern -- Journaling --
The Write Concern w Option -- Write Concern and Secondary Reads --
MongoDB Atlas -- Atlas Search -- Atlas Data Lake -- Summary --
Chapter 14: Sharding -- Sharding Fundamentals -- Scaling
and Sharding -- Sharding Concepts -- To Shard or Not to Shard? --
Shard Key Selection -- Range- vs. Hash-Based Sharding -- Zone
Sharding -- Shard Balance -- Rebalancing Shards -- Modifying
the Balancer Window -- Disabling the Balancer -- Changing the Chunk
Size -- Changing Shard Keys -- Sharded Queries -- Sharded Explain
Plans -- Shard Key Lookups -- Accidental Shard Merge -- Shard Key
Range -- Sorting -- Non-Shard Key Lookups -- Aggregations and Sorts
-- Sharded lookup Operations -- Summary -- Index.
