

1. Record Nr.	UNINA9910494586303321
Autore	Tiepolo Gianluca
Titolo	Getting started with RethinkDB : absorb the knowledge required to utilize, manage, and deploy RethinkDB using Node.js // Gianluca Tiepolo
Pubbl/distr/stampa	Birmingham : , : Packt Publishing, , 2016
ISBN	1-78588-446-8
Edizione	[1st edition]
Descrizione fisica	1 online resource (176 p.)
Collana	Community experience distilled
Soggetti	Internet programming Electronic books.
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Includes index.
Nota di contenuto	Cover ; Copyright; Credits; About the Author; Acknowledgement; About the Reviewer; www.PacktPub.com; Table of Contents; Preface; Chapter 1: Introducing RethinkDB; Rethinking the database; Changefeeds; Horizontal scalability; Powerful query language; Developer-oriented; Document-oriented; Lock-free architecture; Immediate consistency; Secondary indexes; Distributed joins; Installing RethinkDB; Installing RethinkDB on Ubuntu/Debian Linux; Installing RethinkDB on CentOS and Fedora; Installing RethinkDB on OS X; Installing RethinkDB using Homebrew; Building RethinkDB from source Configuring RethinkDBRunning as a daemon; Creating a configuration file; Starting RethinkDB; Running a query; Summary; Chapter 2: The ReQL Query Language; Documents; Document databases; JSON document format; Keys; Arrays; Embedded documents; Data modeling; Introducing ReQL; An explicit query language; Building a query; Inserting data; Batch inserts; Reading data; Filtering results; Manipulating results; Updating data; Updating existing attributes; Adding new attributes; Deleting data; Removing all documents; Deleting a table; Deleting a database; Summary Chapter 3: Clustering, Sharding, and ReplicationAn introduction to scaling; What kind of system is it?; Scaling reads; Scaling writes; Scaling data; Clustering RethinkDB; Creating a cluster; Adding a server to the cluster; Running queries on the cluster; Replication; Adding a

secondary replica; Failover; Sharding; Sharding a table; Summary; Chapter 4: Performance Tuning and Advanced Queries; Performance tuning; Increasing the cache size; Increasing concurrency; Using soft durability mode; Bulk data import; Introducing indexing; Evaluating query performance; Creating and using an index Compound indexesAdvanced queries; Limits, skips, and sorts; The limit command; The skip command; Sorting documents; Finding a random document; Grouping; Aggregations; Average; Maximum; The pluck command; Summary; Chapter 5: Programming RethinkDB in Node.js; Introducing Node.js; An increasingly popular technology; An event-driven design; Installing Node.js; Installing on Linux; Installing on Mac OS X; Running Node.js; Installing the RethinkDB module; Connecting to RethinkDB; Running a simple query; Inserting documents; Reading documents; Updating and deleting documents Introducing ChangefeedsA simple example using Changefeeds; Summary; Chapter 6: RethinkDB Administration and Deployment; RethinkDB administration tools; Backing up your data; Backing up a single table; Setting up automatic backups; Restoring your data; Securing RethinkDB; Securing the web interface; Securing the driver port; Monitoring RethinkDB; Monitoring issues; Monitoring running jobs; Deploying RethinkDB; Summary; Chapter 7: Developing Real-Time Web Applications; Introducing real-time web applications; Examples of real-time web apps; Going real time on the Web; Polling; AJAX; WebSockets Developing web applications with Node.js

Sommario/riassunto

Absorb the knowledge required to utilize, manage, and deploy RethinkDB using Node.js About This Book Make the most of this open source, scalable database - RethinkDB - to ease the construction of web applications Run powerful queries using ReQL, which is the most convenient language to manipulate JSON documents with Develop fully-fledged real-time web apps using Node.js and RethinkDB Who This Book Is For Getting Started with RethinkDB is ideal for developers who are new to RethinkDB and need a practical understanding to start working with it. No previous knowledge of database programming is required, although a basic knowledge of JavaScript or Node.js would be helpful. What You Will Learn Download and install the database on your system Configure RethinkDB's settings and start using the web interface Import data into RethinkDB Run queries using the ReQL language Create shards, replicas, and RethinkDB clusters Use an index to improve database performance Get to know all the RethinkDB deployment techniques In Detail RethinkDB is a high-performance document-oriented database with a unique set of features. This increasingly popular NoSQL database is used to develop real-time web applications and, together with Node.js, it can be used to easily deploy them to the cloud with very little difficulty. Getting Started with RethinkDB is designed to get you working with RethinkDB as quickly as possible. Starting with the installation and configuration process, you will learn how to start importing data into the database and run simple queries using the intuitive ReQL query language. After successfully running a few simple queries, you will be introduced to other topics such as clustering and sharding. You will get to know how to set up a cluster of RethinkDB nodes and spread database load across multiple machines. We will then move on to advanced queries and optimization techniques. You will discover how to work with RethinkDB from a Node.js environment and find out all about deployment techniques. Finally, we'll finish by working on a fully-fledged example that uses the Node.js framework and advanced features such as Changefeeds to develop a real-time web application. Style and approach This is a step-by-step

book that provides a practical approach to RethinkDB programming,
and is explained in a conversational, easy-to-follow style.
