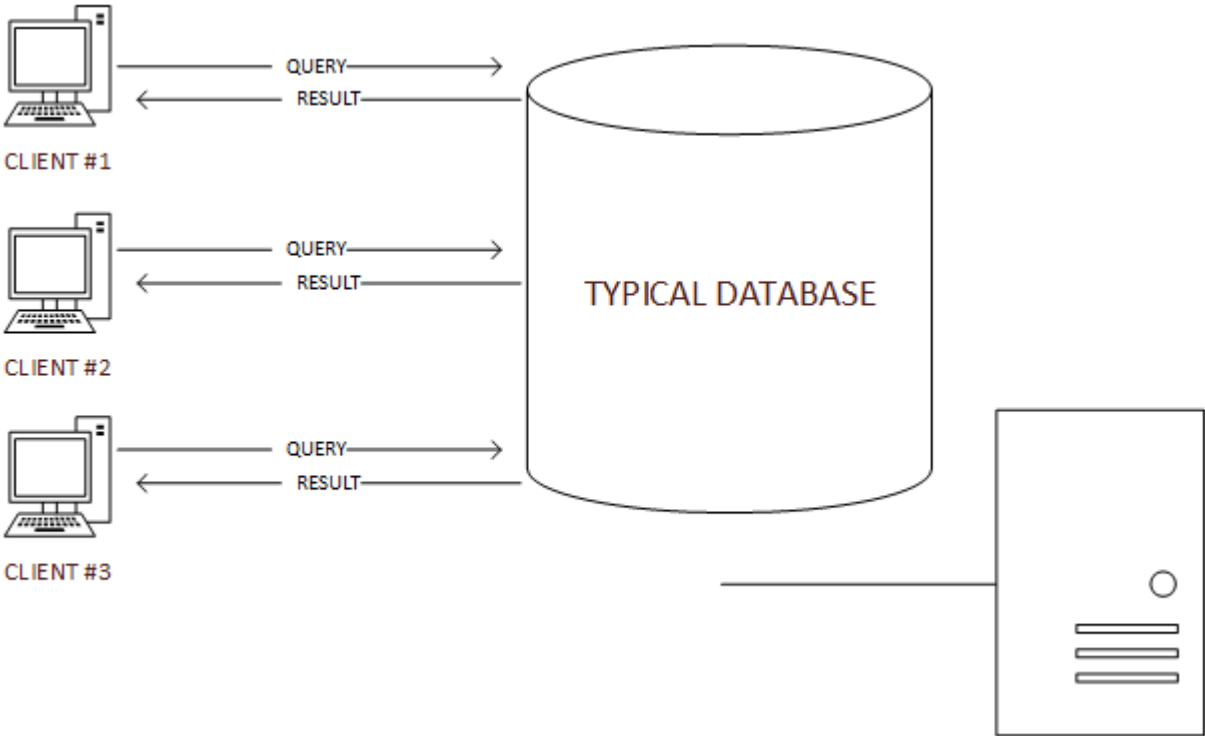
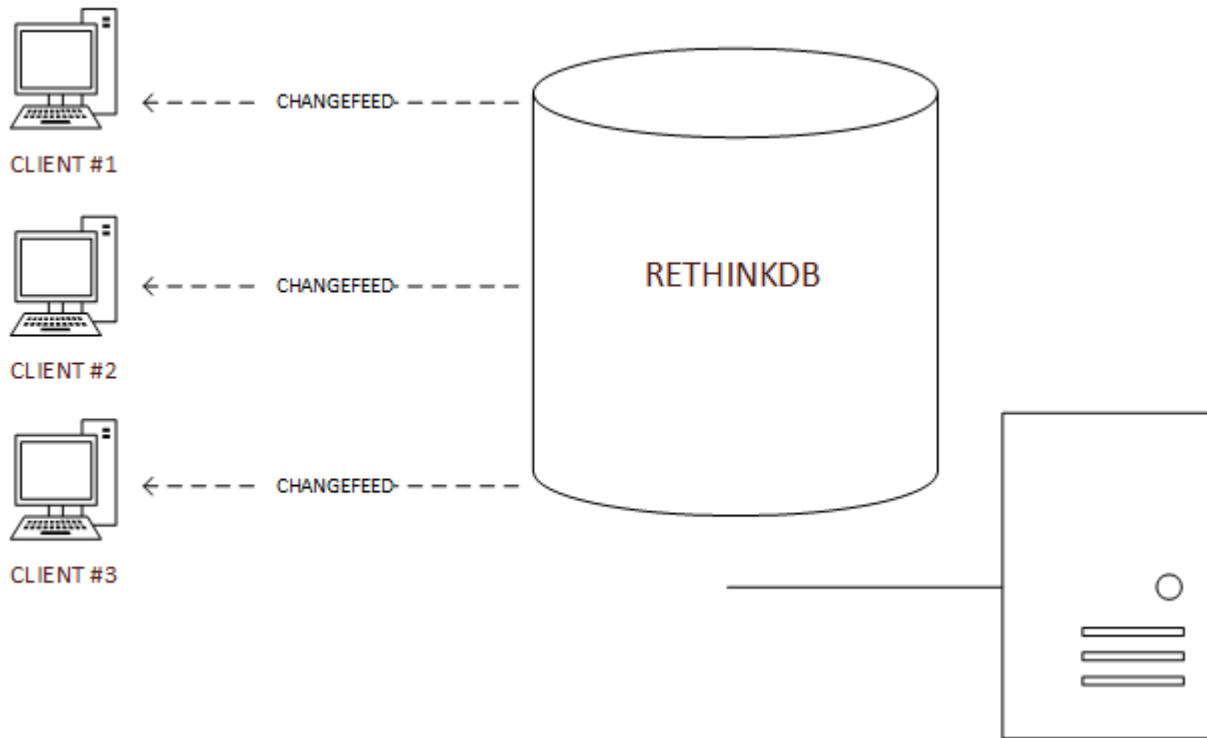


# Chapter 1: Introducing RethinkDB





```

ubuntu@ip-172-31-19-125:~$ sudo apt-get install rethinkdb
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following NEW packages will be installed:
  rethinkdb
0 upgraded, 1 newly installed, 0 to remove and 35 not upgraded.
Need to get 6,201 kB of archives.
After this operation, 24.3 MB of additional disk space will be used.
Get:1 http://download.rethinkdb.com/apt/ trusty/main rethinkdb amd64 2.0.3~0trusty [6,201 kB]
Fetched 6,201 kB in 1s (4,476 kB/s)
Selecting previously unselected package rethinkdb.
(Reading database ... 86567 files and directories currently installed.)
Preparing to unpack ../rethinkdb_2.0.3~0trusty_amd64.deb ...
Unpacking rethinkdb (2.0.3~0trusty) ...
Processing triggers for ureadahead (0.100.0-16) ...
Processing triggers for man-db (2.6.7.1-1ubuntu1) ...
Setting up rethinkdb (2.0.3~0trusty) ...
Adding system startup for /etc/init.d/rethinkdb ...
/etc/rc0.d/K20rethinkdb -> ../init.d/rethinkdb
/etc/rc1.d/K20rethinkdb -> ../init.d/rethinkdb
/etc/rc6.d/K20rethinkdb -> ../init.d/rethinkdb
/etc/rc2.d/S20rethinkdb -> ../init.d/rethinkdb
/etc/rc3.d/S20rethinkdb -> ../init.d/rethinkdb
/etc/rc4.d/S20rethinkdb -> ../init.d/rethinkdb
/etc/rc5.d/S20rethinkdb -> ../init.d/rethinkdb
rethinkdb: No instances defined in /etc/rethinkdb/instances.d/
rethinkdb: See http://www.rethinkdb.com/docs/guides/startup/ for more information
The RethinkDB startup service is installed but disabled. To enable it,
follow the instructions in the guide located at http://www.rethinkdb.com/docs/guides/startup/
Processing triggers for ureadahead (0.100.0-16) ...
ubuntu@ip-172-31-19-125:~$

```



```
[300/321] CC build/release/obj/clustering/immediate_consistency/branch/listener.o
[301/321] CC build/release/obj/clustering/immediate_consistency/branch/backfiller.o
[302/321] CC build/release/obj/clustering/immediate_consistency/branch/replier.o
[303/321] CC build/release/obj/clustering/immediate_consistency/branch/history.o
[304/321] CC build/release/obj/clustering/immediate_consistency/branch/metadata.o
[305/321] CC build/release/obj/clustering/immediate_consistency/branch/broadcaster.o
[306/321] CC build/release/obj/clustering/immediate_consistency/query/master_access.o
[307/321] CC build/release/obj/clustering/immediate_consistency/query/direct_reader.o
[308/321] CC build/release/obj/clustering/immediate_consistency/query/master.o
[309/321] CC build/release/obj/clustering/immediate_consistency/query/metadata.o
[310/321] CC build/release/obj/clustering/generic/multi_client_client.o
[311/321] CC build/release/obj/clustering/generic/nonoverlapping_regions.o
[312/321] BUILD v8_3.30.33.16
[313/321] BUILD re2_20140111
[314/321] LD build/release/rethinkdb
[315/321] CP build/release/rethinkdb-gdb.py
[316/321] CP ./build/drivers/ruby/lib
[317/321] CONVERT_PROTOFILE build/drivers/ruby/lib/ql2.pb.rb
[318/321] CONVERT_PROTOFILE drivers/ruby/lib/ql2.pb.rb
[319/321] CP ./build/drivers/python/rethinkdb/
[320/321] CONVERT_PROTOFILE build/drivers/python/rethinkdb/ql2_pb2.py
[321/321] CONVERT_PROTOFILE drivers/python/rethinkdb/ql2_pb2.py
```

**RethinkDB** Dashboard Tables Servers Data Explorer Logs

Connected to **rethinkdb1** Issues **No issues** Servers **1/1 connected** Tables **0/0 ready**

<b>Availability</b> 0 issues detected 0/0 primaries available	<b>Redundancy</b> 0 issues detected 0/0 replicas available	<b>Connectivity</b> 0 issues detected 1/1 servers connected	<b>Indexes</b> 0 issues detected 0 indexes constructing
---	--	---	---

**Cluster performance** ● Reads/sec: 0 ● Writes/sec: 0

20K  
15K  
10K  
5K  
0

04:26:30 04:26:45 04:27:00 04:27:15 04:27:30 04:27:45 04:28:00 04:28:15

**Tables in the cluster** + Add Database - Delete selected tables

**DATABASE** **test** + Add Table Delete Database

There are no tables in this database.

**Data Explorer** History Settings Fullscreen

```
1 r.db('test').tableCreate('people')
```

Clear Run

## Chapter 2: The ReQL Query Language

**Data Explorer** History Settings Fullscreen

```
1 r.db('test').table('people').insert({
2   "name": "Alex",
3   "surname": "Jones",
4   "yearOfBirth": 1990,
5   "email": "alexjones@gmail.com",
6   "address": {
7     "street": "23, St. John's Street",
8     "city": "London",
9     "postcode": "EC1V 4PY"
10  },
11  "interests": {
12    "programming",
13    "football",
14    "cars"
15  }
16 }
17 )
```

Clear Run

Executed in 141ms. 1 row returned

Tree view Table view Raw view Query profile

```
{
  "deleted": 0,
  "errors": 0,
  "generated_keys": [
    "c4cad21f-edf2-4ae1-8269-73dce5ff8bff"
  ],
  "inserted": 1,
  "replaced": 0,
  "skipped": 0,
  "unchanged": 0
}
```

Executed in 134ms. 1 row returned

Tree view Table view Raw view Query profile

```
{
  "address": {
    "city": "London",
    "postcode": "EC1V 4PY",
    "street": "23, St. John's Street"
  },
  "email": "alexjones@gmail.com",
  "id": "c4cad21f-edf2-4ae1-8269-73dce5ff8bff",
  "interests": [
    "programming",
    "football",
    "cars"
  ],
  "name": "Alex",
  "surname": "Jones",
  "yearOfBirth": 1990
}
```

## Data Explorer

History



```
1 r.db('test').table('people').insert([{\n2   "name": "Ted",\n3   "surname": "Wilson",\n4   "yearOfBirth": 1985,\n5   "email": "ted_wilson@gmail.com",\n6   "address": {\n7     "street": "144, Oxford Street",\n8     "city": "London",\n9     "postcode": "W1D 1LZ"\n10  },\n11  "interests": [\n12    "programming",\n13    "big data"\n14  ],\n15  }, {\n16  "name": "Amy",\n17  "surname": "Alcott",\n18  "yearOfBirth": 1998,\n19  "email": "amy_98@gmail.com",\n20  "address": {\n21    "street": "98, Sutton New Rd",\n22    "city": "Birmingham",\n23    "postcode": "B23 6RE"\n24  },\n25  "interests": [\n26    "dance",\n27    "fashion",\n28    "shopping"\n29  ]\n30  },\n31  ]\n32  })
```

## Data Explorer

History



```
1 r.db('test').table('people')
```

Clear

Run

Executed in 119ms. 3 rows returned

Tree view

Table view

Raw view

Query profile

	address.city	address.postcode	address.street	email	id
1	London	W1D 1LZ	144, Oxford Street	ted_wilson@gmail.com	083d0878-808b-4ebf-abee-8e95
2	London	EC1V 4PY	23, St. John's Street	alexjones@gmail.com	c4cad21f-edf2-4ae1-8269-73dc
3	Birmingham	B23 6RE	98, Sutton New Rd	amy_98@gmail.com	f1664276-1aad-4998-8240-410e

Executed in 147ms. 2 rows returned

Tree view

Table view

Raw view

Query profile

	address.city	address.postcode	address.street	email	id
1	London	W1D 1LZ	144, Oxford Street	ted_wilson@gmail.com	083d0878-808b-4ebf-abee-8e95
2	London	EC1V 4PY	23, St. John's Street	alexjones@gmail.com	c4cad21f-edf2-4ae1-8269-73dc

Executed in 116ms. 2 rows returned

Tree view

Table view

Raw view

Query profile

street	email	id	interests	name	surname	yearOfBirth
in's Street	alexjones@gmail.com	c4cad21f-edf2-4ae1-8269-73dce5ff8bff	▶ [ ... ]	Alex	Jones	1990
New Rd	amy_98@gmail.com	f1664276-1aad-4998-8240-410a82883115	▶ [ ... ]	Amy	Alcott	1998

## Data Explorer

History



```
1 r.table('people').filter(r.row("yearOfBirth").ge(1990)).pluck("name", "email", "yearOfBirth")
```

Clear

Run

Executed in 142ms. 2 rows returned

Tree view

Table view

Raw view

Query profile

```
{
  "email": alexjones@gmail.com, »
  "name": "Alex" ,
  "yearOfBirth": 1990
}
{
  "email": amy_98@gmail.com, »
  "name": "Amy" ,
  "yearOfBirth": 1998
}
```

## Data Explorer

History



```
1 r.table('people').filter(r.row("name").eq("Amy")).pluck("name", "yearOfBirth")
```

Clear

Run

Executed in 116ms. 1 row returned

Tree view

Table view

Raw view

Query profile

```
{
  "name": "Amy" ,
  "yearOfBirth": 1997
}
```

## Data Explorer

History



```
1 r.table('people').pluck("name", "nationality")
```

Clear

Run

Executed in 123ms. 3 rows returned

Tree view

Table view

Raw view

Query profile

	name	nationality
1	Ted	English
2	Alex	Spanish
3	Amy	English

## Data Explorer

History



```
1 r.table('people').count()
```

Clear

Run

Executed in 138ms. 1 row returned

Tree view

Table view



Raw view

Query profile





0






## Chapter 3: Clustering, Sharding, and Replication



**RethinkDB** Dashboard Tables Servers Data Explorer Logs

 Connected to **rethink1**  Issues **No issues**  Servers **1/1 connected**  Tables **0/0 ready**

**Servers connected to the cluster**

 <b>rethink1</b> default	 0 primaries, 0 secondaries	 Connected
---	--	---



Connected to  
**rethink1**



Issues  
**No issues**



Servers  
**2/2 connected**



Tables  
**0/0 ready**

● **Availability**

0 issues detected  
0/0 primaries available

● **Redundancy**

0 issues detected  
0/0 replicas available

● **Connectivity**

0 issues detected  
2/2 servers connected

● **Indexes**

0 issues detected  
0 indexes constructing

**Cluster performance**

● Reads/sec: 0   ● Writes/sec: 0



## Data Explorer

History



```
1 r.db('test').tableCreate('clusteringTest1')
```

Clear

Run

1 row returned in 5.02s

Tree view

Table view

Raw view

Query profile

```
{
  "config_changes": [
    {
      "new_val": {
        "db": "test",
        "durability": "hard",
        "id": "b5d2f06e-943d-46af-9d78-72c533e799e6",
        "name": "clusteringTest1",
        "primary_key": "id",
        "shards": [
          {
            "primary_replica": "rethink2",
            "replicas": [
              "rethink2"
            ]
          }
        ],
        "write_acks": "majority"
      },
      "old_val": null
    }
  ],
  "tables_created": 1
}
```

1 row returned in 917ms.

Tree view

Table view

Raw view

Query profile

```
{
  "config_changes": [
    {
      "new_val": {
        "db": "test",
        "durability": "hard",
        "id": "2695dd3e-6cb4-4893-9437-7b3762cd6037",
        "name": "clusteringTest2",
        "primary_key": "id",
        "shards": [
          {
            "primary_replica": "rethink1",
            "replicas": [
              "rethink1"
            ]
          }
        ],
        "write_acks": "majority"
      },
      "old_val": null
    }
  ],
  "tables_created": 1
}
```

RethinkDB

Dashboard

Tables

Servers

Data Explorer

Logs



Search



Connected to  
**rethink1**



Issues  
**No issues**



Servers  
**2/2 connected**



Tables  
**2/2 ready**

### Servers connected to the cluster

<b>rethink2</b> default	1 primary, 0 secondaries	Connected
<b>rethink1</b> default	1 primary, 0 secondaries	Connected

## Data Explorer

History



```
1 r.db('test').tableCreate('people')
```

Clear

Run

1 row returned in <1ms.

Tree view

Table view

Raw view

Query profile

```
{
  "config_changes": [
    {
      "new_val": {
        "db": "test",
        "durability": "hard",
        "id": "30a9e82a-1cfa-492a-be4e-b3f2a9d5e38f",
        "name": "people",
        "primary_key": "id",
        "shards": [
          {
            "primary_replica": "rethink2",
            "replicas": [
              "rethink2"
            ]
          }
        ],
        "write_acks": "majority"
      },
      "old_val": null
    }
  ],
  "tables_created": 1
}
```

## Sharding and replication

1 shard

1 replica per shard

Reconfigure

## Sharding and replication for test.people

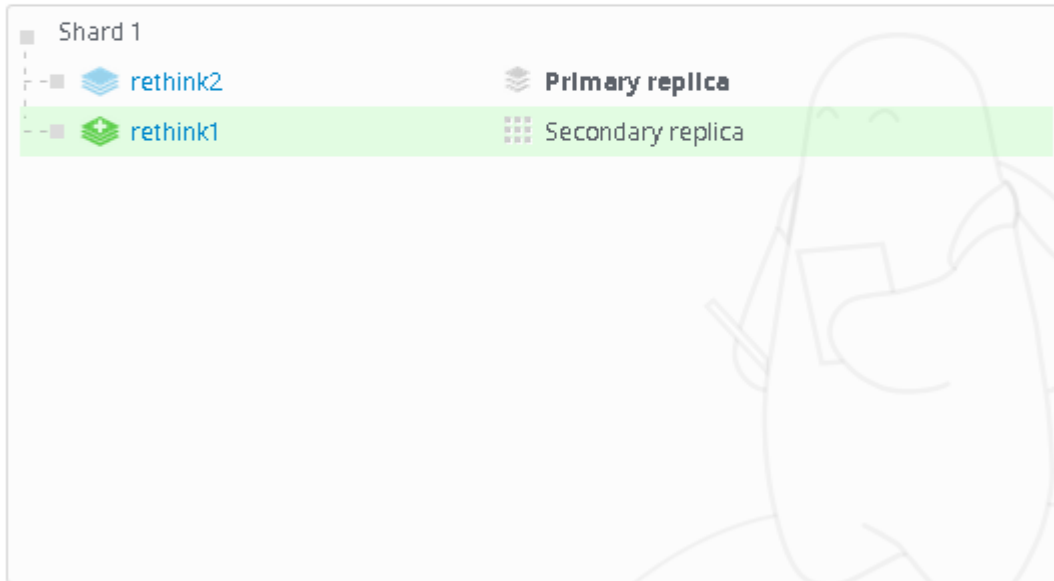


**i** Applying changes may cause the table to become temporarily unavailable.

shards  
max: 2

replicas per shard  
max: 2

▼ Where's my data going?



Cancel

Apply configuration

### Tables in the cluster

+ Add Database

- Delete selected tables

DATABASE test

+ Add Table

Delete Database

rethink people

1 shard, 2 replicas

Ready 2/2

## Data Explorer

History



```
1 r.db('test').table('people').insert({'name': 'Alice', 'age': 23})
```

Clear

Run

1 row returned in 1.65s.

Tree view

Table view

Raw view

Query profile

```
{
  "deleted": 0 ,
  "errors": 0 ,
  "generated_keys": [
    "e0edbb17-a2ca-437e-8614-6f1b139f0a9b"
  ] ,
  "inserted": 1 ,
  "replaced": 0 ,
  "skipped": 0 ,
  "unchanged": 0
}
```

RethinkDB

Dashboard

Tables

Servers

Data Explorer

Logs



Search

Connected to  
rethink1

Issues  
1 issue

Servers  
1/2 connected

Tables  
0/1 ready

1 issue needs to be resolved

Resolve issues

### ● Availability

1/1 tables are affected  
1/1 primaries unavailable

### ● Redundancy

1/1 tables are affected  
1/2 replicas are unavailable

### ● Connectivity

1 issue detected  
1/2 servers are disconnected

### ● Indexes

0 issues detected  
0 indexes constructing

## Data Explorer

History  

```
1 r.db('test').table('people', {'useOutdated': true})
```

Clear Run

5 rows returned in 906ms.

[Tree view](#) [Table view](#) [Raw view](#) [Query profile](#)

	age	id	name
1	31	eedb2032-0253-4c9f-a285-f194eb47e3e3	Bob
2	23	e0edbb17-a2ca-437e-8614-6f1b139f0a9b	Alice
3	60	5080aa43-e112-442a-b8af-862dbed37978	George
4	18	c910a87a-863a-429d-934b-5c170b9bf052	Emily
5	13	ea559eb4-f0fb-46b7-a82a-8c804ed6f660	Amy

## Sharding and replication

1 shard

2 replicas per shard

Reconfigure



# Sharding and replication for test.people

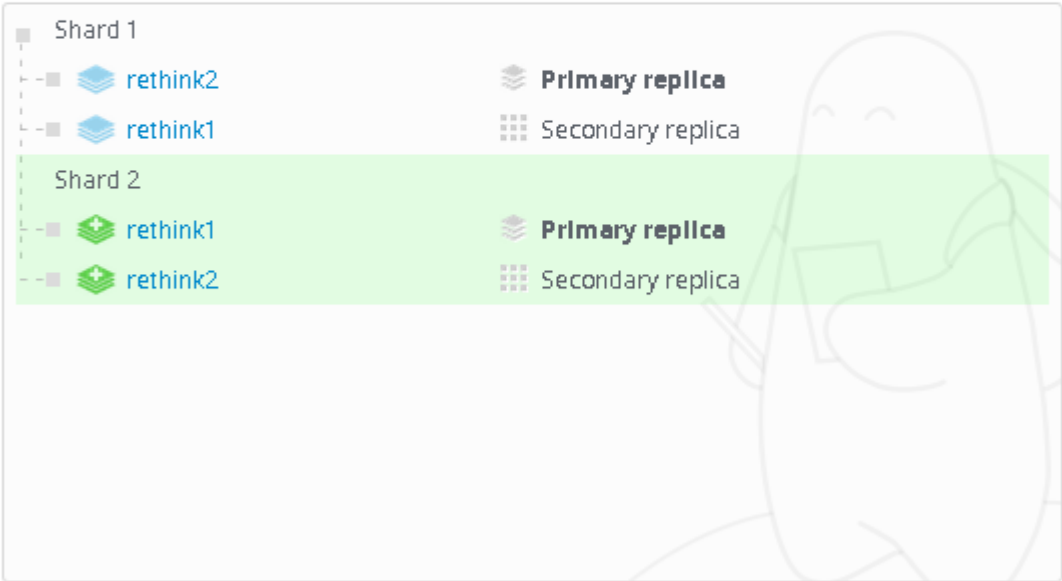


**i** Applying changes may cause the table to become temporarily unavailable.

shards  
max: 2

replicas per shard  
max: 2


▼ Where's my data going?



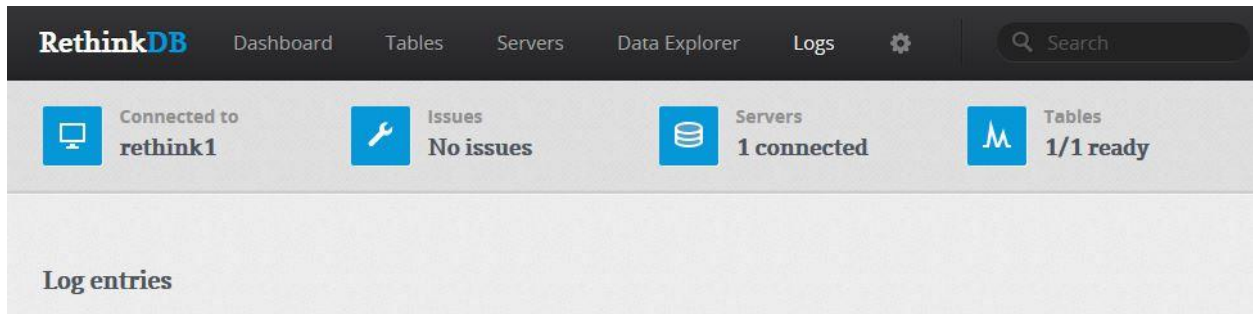
Cancel

Apply configuration

## Data distribution



## Chapter 4: Performance Tuning and Advanced Queries



The dashboard shows the RethinkDB interface with navigation tabs for Dashboard, Tables, Servers, Data Explorer, and Logs. A search bar is located in the top right. Below the navigation, there are four status cards: 'Connected to rethink1', 'Issues: No issues', 'Servers: 1 connected', and 'Tables: 1/1 ready'. A 'Log entries' section is visible below these cards.



Automatically using cache size of 2830 MB  
Posted by rethink1

### Data Explorer



```
1 r.db('test').table('people').count()
```



1 row returned in 20ms.



30000

### Data Explorer



Enable automatic suggestions and command descriptions

Auto-insert pair brackets, parentheses, and quotes

Enable query profiler

Max results per query (Note: returning more results will increase loading time)

## Data Explorer

History



```
1 r.db('test').table('people').filter(r.row('name').eq('Alex'))
```

Clear

Run

40 rows returned. Displaying rows 1-40, more available

Tree view

Table view

Raw view

Query profile

↻ 306ms round-trip time

🕒 132ms server time

📄 8 shard accesses

### Secondary indexes

■ No secondary indexes found.

■


Cancel

Create

### Secondary indexes

You successfully created the secondary index **name**. ✕

■ **name** — delete



Processing

[Create a new secondary index →](#)

## Secondary indexes

▪ **name** — delete

Ready to use

Create a new secondary index →

### Data Explorer

History



```
1 r.db('test').table('people').getAll("Alex", {index: "name"})
```

Clear

Run

40 rows returned. Displaying rows 1-40, more available

Tree view

Table view

Raw view

Query profile

↻ 232ms round-trip time

🕒 1ms server time

📄 8 shard accesses

4 rows returned.

Tree view

Table view

Raw view

Query profile

	age	email	id	name	surname
1	22	judith@crane.bg	815da91f-e921-4cf2-a514-e4ef4fa6da01	Alex	Dyer
2	22	jerome@lynn.ai	90674bef-7796-4a7a-9b60-314634868b61	Alex	Rosenthal
3	22	guy@olsen.bh	fd667db0-e481-4c4a-a731-c01cd5543f5a	Alex	Bowles
4	22	leah@wolf.nl	97313358-819f-4567-bd94-48b967fce076	Alex	Snow

## Data Explorer

History Settings Full Screen

```
1 r.db("test").table("people").getAll(["Alex", 22], {index: "name_and_age"})
```

Clear Run

4 rows returned.

Tree view Table view Raw view Query profile

↻ 149ms round-trip time ⌚ <1ms server time 📄 8 shard accesses

## Data Explorer

History Settings Full Screen

```
1 r.db("test").table("people").orderBy(r.desc('surname')).skip(5).limit(10)
```

Clear Run

10 rows returned in 3.76s.

Tree view Table view Raw view Query profile

↻ 3.87s round-trip time ⌚ 3.76s server time 📄 24 shard accesses

## Data Explorer

History Settings Full Screen

```
1 r.db("test").table("people").orderBy({index: r.desc("surname")}).skip(5).limit(10)
```

Clear Run

10 rows returned. Displaying rows 1-10

Tree view Table view Raw view Query profile

↻ 152ms round-trip time ⌚ 16ms server time 📄 0 shard accesses

438 rows returned in 145ms. Displaying rows 1-438

[Tree view](#)

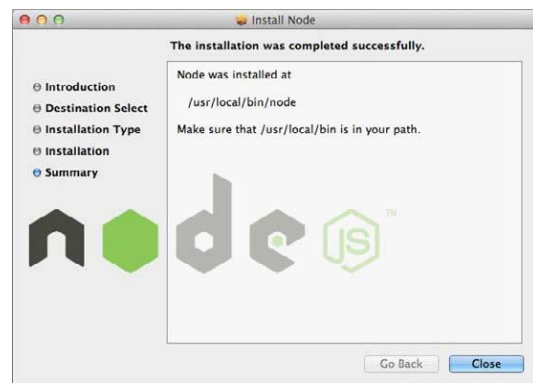
[Table view](#)

[Raw view](#)

[Query profile](#)

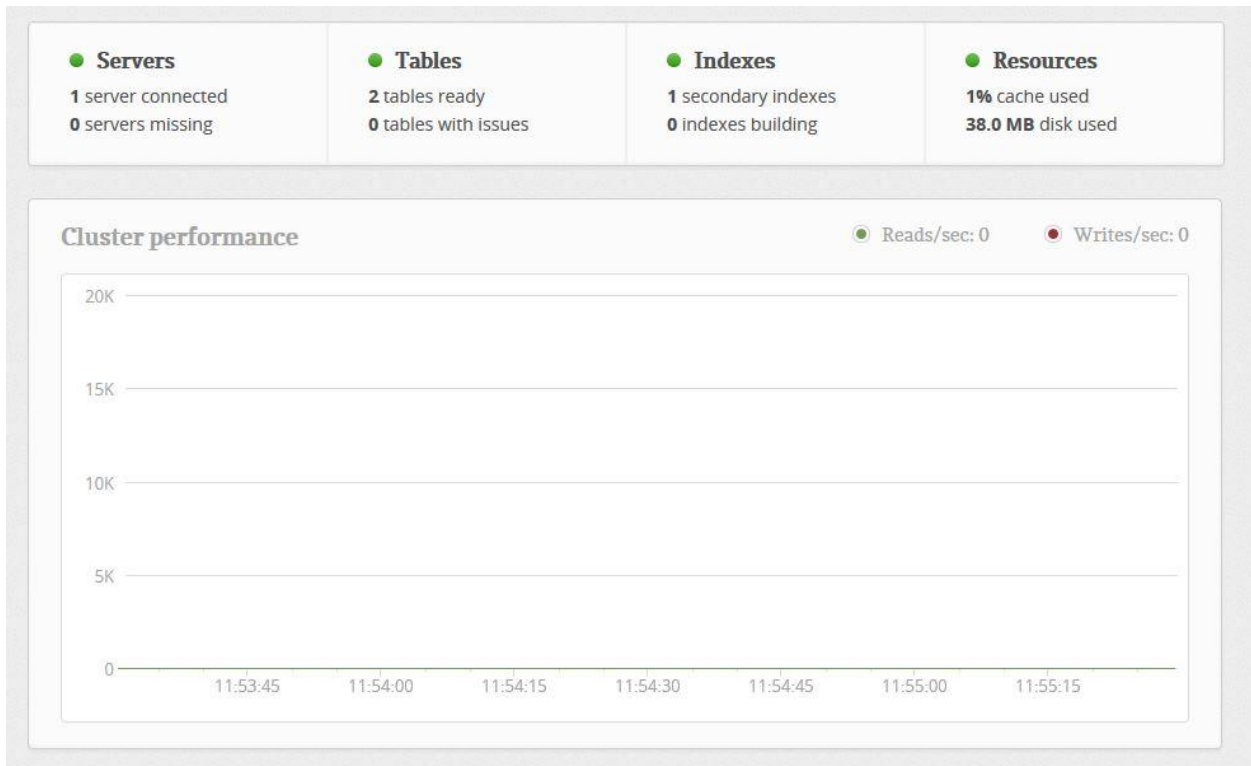
	group	reduction
1	Aaron	40
2	Adam	25
3	Alan	55
4	Albert	26
5	Alex	91
6	Alexander	27
7	Alfred	100
8	Alice	120
9	Alicia	21
10	Alison	51

## Chapter 5: Programming RethinkDB in Node.js





## Chapter 6: RethinkDB Administration and Deployment



### Edit inbound rules ✕

Type <small>i</small>	Protocol <small>i</small>	Port Range <small>i</small>	Source <small>i</small>
Custom TCP Rule ▾	TCP	8080	Custom IP ▾ 127.0.0.1/32 ✕

Add Rule Cancel Save

Quick Start

My AMIs

**AWS Marketplace**

Community AMIs

▼ **Categories**

All Categories

  Software Infrastructure (2)

▼ **Operating System**

  Clear Filter

    ▼ **All Linux/Unix**

Ubuntu (2)

▼ **Software Pricing Plans**

Free (2)

▼ **Region**

Current Region (2)

All Regions (2)

RethinkDB

Free tier eligible

**RethinkDB**

★★★★★ (2) | 1.16.2 [Previous versions](#) | Sold by [RethinkDB](#)

**\$0.00/hr for software + AWS usage fees**

Linux/Unix, Ubuntu 12.04 | 64-bit Amazon Machine Image (AMI) | Updated: 2/26/15

RethinkDB is a scalable, open-source, distributed database system built to store JSON documents. It features a query language that has useful queries like table joins and ...

[More info](#)

RethinkDB

Free tier eligible

**RethinkDB**

★★★★★ (0) | 2.1.5 [Previous versions](#) | Sold by [RethinkDB](#)

**\$0.00/hr for software + AWS usage fees**

Linux/Unix, Ubuntu 14.04 | 64-bit Amazon Machine Image (AMI) | Updated: 10/9/15

RethinkDB is the first open-source, scalable JSON database built from the ground up for the realtime web. It inverts the traditional database architecture by exposing an ...

[More info](#)

### Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. [Learn more about instance types and how they can meet your computing needs.](#)

Filter by: General purpose Current generation Show/Hide Columns

**Currently selected:** t2.small (Variable ECU, 1 vCPU, 2.5 GHz, Intel Xeon Family, 2 GB memory, EBS only)

**Note:** The vendor recommends using a **t2.small** instance (or larger) for the best experience with this product.

	Family	Type	vCPUs (1)	Memory (GB)	Instance Storage (GB) (1)	EBS-Optimized Available (1)	Network Performance (1)
<input type="checkbox"/>	General purpose	t2.micro	1	1	EBS only	-	Low to Moderate
<input checked="" type="checkbox"/>	General purpose	t2.small	1	2	EBS only	-	Low to Moderate
<input type="checkbox"/>	General purpose	t2.medium	2	4	EBS only	-	Low to Moderate
<input type="checkbox"/>	General purpose	t2.large	2	8	EBS only	-	Low to Moderate
<input type="checkbox"/>	General purpose	m4.large	2	8	EBS only	Yes	Moderate
<input type="checkbox"/>	General purpose	m4.xlarge	4	16	EBS only	Yes	High
<input type="checkbox"/>	General purpose	m4.2xlarge	8	32	EBS only	Yes	High
<input type="checkbox"/>	General purpose	m4.4xlarge	16	64	EBS only	Yes	High
<input type="checkbox"/>	General purpose	m4.10xlarge	40	160	EBS only	Yes	10 Gigabit
<input type="checkbox"/>	General purpose	m3.medium	1	3.75	1 x 4 (SSD)	-	Moderate
<input type="checkbox"/>	General purpose	m3.large	2	7.5	1 x 32 (SSD)	-	Moderate
<input type="checkbox"/>	General purpose	m3.xlarge	4	15	2 x 40 (SSD)	Yes	High
<input type="checkbox"/>	General purpose	m3.2xlarge	8	30	2 x 80 (SSD)	Yes	High

Cancel Previous Review and Launch Next: Configure Instance Details



## Initiating Instance Launches

Please do not close your browser while this is loading

Creating security groups... Successful

Authorizing inbound rules... Successful

**Subscribing to Product...**

Chapter 7: Developing Real-Time Web Applications

