



COUCHDB

nosql document store



tutorialspoint

SIMPLY EASY LEARNING

www.tutorialspoint.com

 <https://www.facebook.com/tutorialspointindia>

 <https://twitter.com/tutorialspoint>

About the Tutorial

This tutorial provides a brief knowledge about CouchDB, the procedures to set it up, and the ways to interact with CouchDB server using cURL and Futon. It also tells how to create, update and delete databases and documents.

Audience

This tutorial helps the professionals aspiring to make a career in Big Data and NoSQL databases, especially the documents store.

Prerequisites

Before you start proceeding with this tutorial, we are assuming that you have a brief knowledge on Big Data, Hadoop, and also have the basic understanding of databases.

Copyright & Disclaimer

© Copyright 2015 by Tutorials Point (I) Pvt. Ltd.

All the content and graphics published in this e-book are the property of Tutorials Point (I) Pvt. Ltd. The user of this e-book is prohibited to reuse, retain, copy, distribute or republish any contents or a part of contents of this e-book in any manner without written consent of the publisher.

We strive to update the contents of our website and tutorials as timely and as precisely as possible, however, the contents may contain inaccuracies or errors. Tutorials Point (I) Pvt. Ltd. provides no guarantee regarding the accuracy, timeliness or completeness of our website or its contents including this tutorial. If you discover any errors on our website or in this tutorial, please notify us at contact@tutorialspoint.com

Table of Contents

About the Tutorial	
Audience.....	i
Prerequisites.....	i
Copyright & Disclaimer	i
Table of Contents.....	ii
1. COUCHDB – INTRODUCTION.....	1
RDBMS.....	1
OLAP	1
NoSQL Databases.....	1
What is CouchDB?.....	2
Why CouchDB?	2
Data Model.....	2
Features of CouchDB: Reduce the Content.....	3
History	4
2. COUCHDB – INSTALLATION.....	5
Installing CouchDB in Windows.....	5
Installing CouchDB in Linux Systems	7
3. COUCHDB – CURL AND FUTON	12
cURL Utility	12
Using cURL Utility	12
cURL Utility Options.....	12
Hello CouchDB	16
List of All Databases	17
Creating a Database	17

Getting Database Info	18
Futon	19
Starting Futon	19
4. COUCHDB – HTTP API	21
HTTP Request Formats	21
HTTP Request Headers	21
Response Headers.....	22
Status Codes	22
HTTP URL Paths.....	23
5. COUCHDB – CREATING A DATABASE.....	25
Creating a Database using cURL Utility.....	25
Creating a Database using Futon.....	26
6. COUCHDB – DELETING A DATABASE	28
Deleting a Database using cURL Utility.....	28
Deleting a Database using Futon.....	29
7. COUCHDB – CREATING A DOCUMENT	31
Creating a Document using cURL Utility	31
Creating a Document using Futon	32
8. COUCHDB – UPDATING A DOCUMENT.....	36
Updating Documents using cURL	36
Updating Documents using Futon	37
9. COUCHDB – DELETING A DOCUMENT.....	40
Deleting a Document using cURL Utility	40
Deleting a Document using Futon	41

10. COUCHDB – ATTACHING FILES.....	43
Attaching Files using cURL.....	43
Attaching Files using Futon	44

1. COUCHDB – INTRODUCTION

Database management system provides mechanism for storage and retrieval of data. There are three main types of database management systems namely RDBMS (Relational Database management Systems), OLAP (Online Analytical Processing Systems) and NoSQL.

RDBMS

RDBMS stands for Relational Database Management System. RDBMS is the basis for SQL, and for all modern database systems like MS SQL Server, IBM DB2, Oracle, MySQL, and Microsoft Access.

A Relational database management system (RDBMS) is a database management system (DBMS) that is based on the relational model as introduced by E. F. Codd.

The data in RDBMS is stored in database objects called **tables**. The table is a collection of related data entries and it consists of columns and rows. It stores only structured data.

OLAP

Online Analytical Processing Server (OLAP) is based on the multidimensional data model. It allows managers and analysts to get an insight of the information through fast, consistent, and interactive access to information.

NoSQL Databases

A NoSQL database (sometimes called as Not Only SQL) is a database that provides a mechanism to store and retrieve data other than the tabular relations used in relational databases. These databases are schema-free, support easy replication, have simple API, eventually consistent, and can handle huge amounts of data (big data).

The primary objective of a NoSQL database is to have the following:

- Simplicity of design,
- Horizontal scaling, and
- Finer control over availability.

NoSQL databases use different data structures compared to relational databases. It makes some operations faster in NoSQL. The suitability of a given NoSQL database depends on the problem it must solve. These databases store both structured data and unstructured data like audio files, video files, documents, etc. These NoSQL databases are classified into three types and they are explained below.

Key-value Store: These databases are designed for storing data in key-value pairs and these databases will not have any schema. In these databases, each data value consists of an indexed key and a value for that key.

Examples: BerkeleyDB, Cassandra, DynamoDB, Riak.

Column Store: In these databases, data is stored in cells grouped in columns of data, and these columns are further grouped into Column families. These column families can contain any number of columns.

Examples: BigTable, HBase, and HyperTable.

Document Store: These are the databases developed on the basic idea of key-value stores where “documents” contain more complex data. Here, each document is assigned a unique key, which is used to retrieve the document. These are designed for storing, retrieving, and managing document-oriented information, also known as semi-structured data.

Examples: CouchDB and MongoDB.

What is CouchDB?

CouchDB is an open source database developed by Apache software foundation. The focus is on the ease of use, embracing the web. It is a NoSQL document store database.

It uses Json, to store data (documents), java script as its query language to transform the documents, http protocol for api to access the documents, query the indices with the web browser. It is a multi master application released in 2005 and it became an apache project in 2008.

Why CouchDB?

- CouchDB have an HTTP-based REST API, which helps to communicate with the database easily. And the simple structure of HTTP resources and methods (GET, PUT, DELETE) are easy to understand and use.
- As we store data in the flexible document-based structure, there is no need to worry about the structure of the data.

- Users are provided with powerful data mapping, which allows querying, combining, and filtering the information.
- CouchDB provides easy-to-use replication, using which you can copy, share, and synchronize the data between databases and machines.

Data Model

- Database is the outermost data structure/container in CouchDB.
- Each database is a collection of independent documents.
- Each document maintains its own data and self-contained schema, an application may access multiple databases.
- Document metadata contains revision information, which makes it possible to merge the differences occurred while the databases were disconnected.
- CouchDB implements multi version concurrency control, to avoid the need to lock the d/b field during writes.

Features of CouchDB: Reduce the Content

Document Storage

CouchDB is a document storage NoSQL database. It provides the facility of storing documents with unique names, and it also provides an API called RESTful HTTP API for reading and updating (add, edit, delete) database documents.

In CouchDB, documents are the primary unit of data and they also include metadata. Document fields are uniquely named and contain values of varying types (text, number, Boolean, lists, etc.), and there is no set limit to text size or element count.

Document updates (add, edit, delete) follow Atomicity, i.e., they will be saved completely or not saved at all. The database will not have any partially saved or edited documents.

Json Document Structure

```
{  
  "field" : "value",  
  "field" : "value",  
  "field" : "value",  
}
```



```
}
```

ACID Properties

CouchDB contains ACID properties as one of its features.

Consistency: When the data in CouchDB was once committed, then this data will not be modified or overwritten. Thus, CouchDB ensures that the database file will always be in a consistent state.

A multi-Version Concurrency Control (MVCC) model is used by CouchDB reads, because of which the client will see a consistent snapshot of the database from the beginning to the end of the read operation.

Whenever a documents is updated, CouchDB flushes the data into the disk, and the updated database header is written in two consecutive and identical chunks to make up the first 4k of the file, and then synchronously flushed to disk. Partial updates during the flush will be discarded.

If the failure occurred while committing the header, a surviving copy of the previous identical headers will remain, ensuring coherency of all previously committed data. Except the header area, consistency checks or fix-ups after a crash or a power failure are never necessary.

Compaction

Whenever the space in the database file got wasted above certain extent, all the active data will be copied (cloned) to a new file. When the copying process is entirely done, then the old file will be discarded. All this is done by compaction process. The database remains online during the compaction and all updates and reads are allowed to complete successfully.

Views

Data in CouchDB is stored in semi-structured documents that are flexible with individual implicit structures, but it is a simple document model for data storage and sharing. If we want see our data in many different ways, we need a way to filter, organize and report on data that hasn't been decomposed into tables.

To solve this problem, CouchDB provides a view model. Views are the method of aggregating and reporting on the documents in a database, and are built on demand to aggregate, join and report on database documents. Because views are built

dynamically and don't affect the underlying document, you can have as many different view representations of the same data as you like.

History

- CouchDB was written in Erlang programming language.
- It was started by Damien Katz in 2005.
- CouchDB became an Apache project in 2008.

The current version of CouchDB is 1.61.

2. COUCHDB – INSTALLATION

This chapter teaches you how to install CouchDB in windows as well as Linux systems.

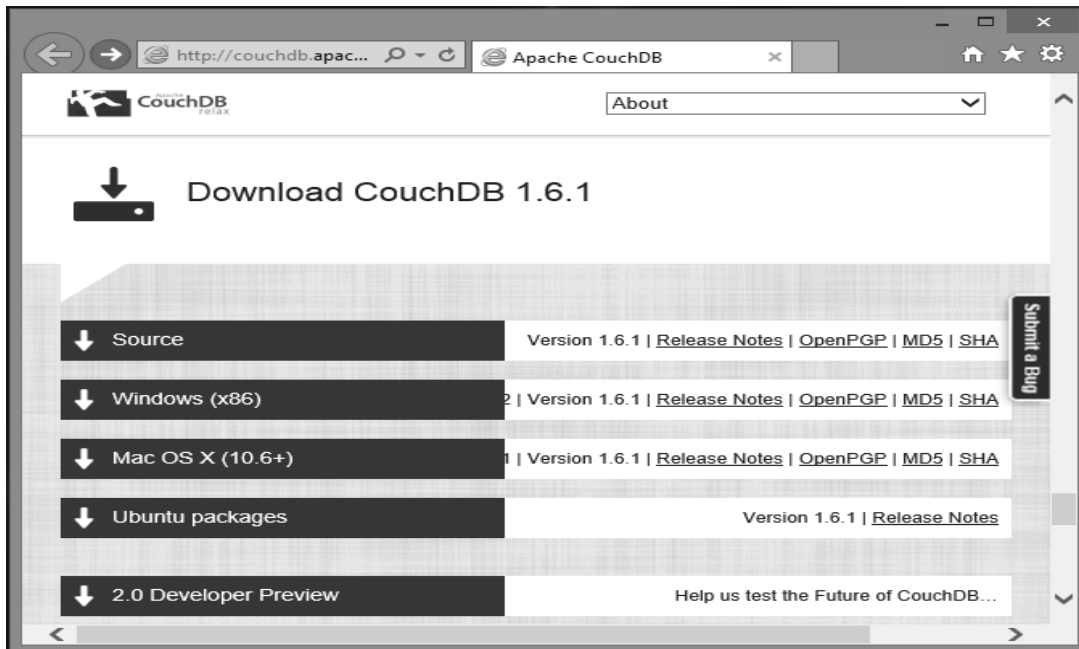
Installing CouchDB in Windows

Download CouchDB

The official website for CouchDB is <http://couchdb.apache.org>. If you click the given link, you can get the home page of the CouchDB official website as shown below.



If you click on the download button that will lead to a page where download links of CouchDB in various formats are provided. The following snapshot illustrates the same.



Choose the download link for windows systems and select one of the provided mirrors to start your download.

Installing CouchDB

CouchDB will be downloaded to your system in the form of setup file named **setup-couchdb-1.6.1_R16B02.exe**. Run the setup file and proceed with the installation.

After installation, open built in web interface of CouchDB by visiting the following link: <http://127.0.0.1:5984/>. If everything goes fine, this will give you a web page, which have the following output.

```
{  
  
  "couchdb":"Welcome","uuid":"c8d48ac61bb497f4692b346e0f400d60",  
  "version":"1.6.1",  
  "vendor":{  
    "version":"1.6.1","name":"The Apache Software Foundation"  
  }  
}
```

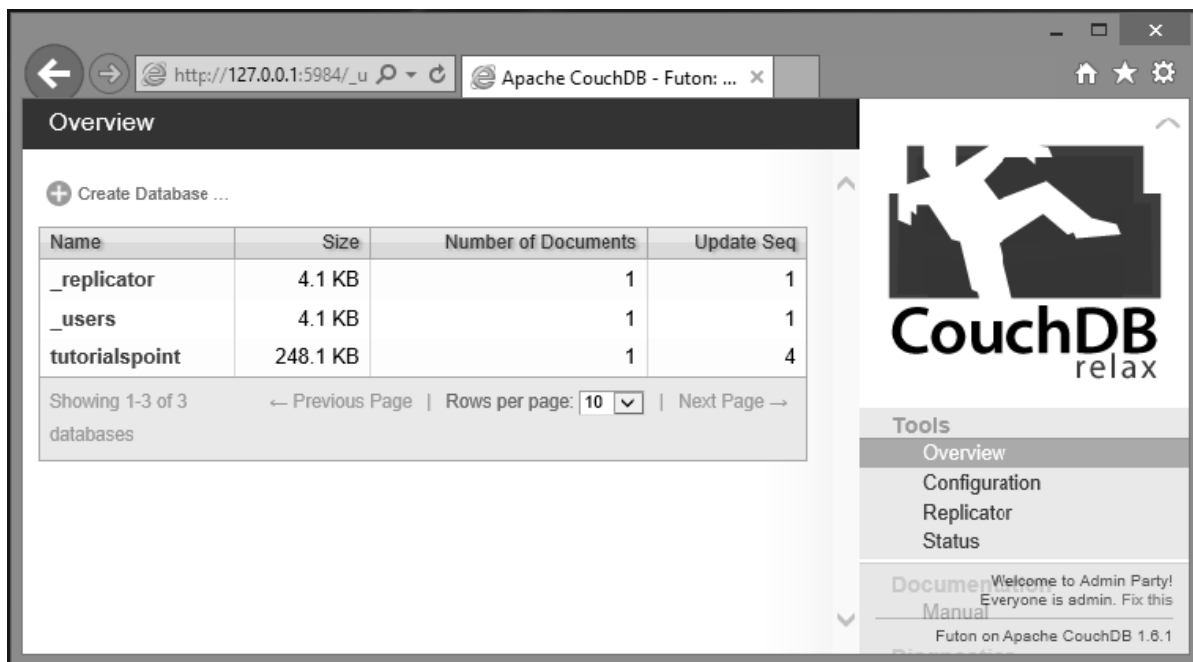
11

```
}
```

You can interact with the CouchDB web interface by using the following url:

```
http://127.0.0.1:5984/_utils/
```

This shows you the index page of Futon, which is the web interface of CouchDB.



Installing CouchDB in Linux Systems

For many of the Linux flavored systems, they provide CouchDB internally. To install this CouchDB follow the instructions.

On Ubuntu and Debian you can use:

```
sudo aptitude install couchdb
```

On Gentoo Linux there is a CouchDB ebuild available:

```
sudo emerge couchdb
```

If your Linux system does not provide CouchDB internally, you have to install CouchDB and its dependencies by your own as shown below.

Installing CouchDB Dependencies

Following is the list of dependencies that are to be installed to get CouchDB in your system.

- Erlang OTP
- ICU
- OpenSSL
- Mozilla SpiderMonkey
- GNU Make
- GNU Compiler Collection
- libcurl
- help2man
- Python for docs
- Python Sphinx

To install these dependencies, type the following commands in the terminal. Here we are using Centos 6.5 and the following commands will install the required softwares compatible to Centos 6.5.

```
$sudo yum install autoconf
$sudo yum install autoconf-archive
$sudo yum install automake
$sudo yum install curl-devel
$sudo yum install erlang-asn1
$sudo yum install erlang-erts
$sudo yum install erlang-eunit
$sudo yum install erlang-os_mon
$sudo yum install erlang-xmerl
$sudo yum install help2man
$sudo yum install js-devel
$sudo yum install libicu-devel
$sudo yum install libtool
$sudo yum install perl-Test-Harness
```

Note: For all these commands you need to use sudo. The following procedure converts a normal user to a sudoer.

- Login as root in admin user.
- Open **sudo** file using the following command:

```
visudo
```

- Then edit as shown below to give your existing user the sudoer privileges:

```
Hadoop All=(All) All , and press esc : x to write the changes to the file.
```

After downloading all the dependencies in your system, download CouchDB following the given instructions.

Downloading CouchDB

Apache software foundation will not provide the complete .tar file for CouchDB, so you have to install it from the source. Download the source file of CouchDB from the following link:

www.apache.org/dist/couchdb/source/1.6.1/apache-couchdb-1.6.1.tar.gz. Create a new directory to install CouchDB, browse to such created directory and download CouchDB source by executing the following commands:

```
$ cd
$ mkdir CouchDB
$ cd CouchDB/
$ wget
http://www.google.com/ur1?q=http%3A%2F%2Fwww.apache.org%2Fdist%2Fcouchdb%2F
source%2F1.6.1%2Fapache-couchdb-1.6.1.tar.gz
```

This will download CouchDB source file into your system. Now unzip the **apache-couchdb-1.6.1.tar.gz** as shown below.

```
$ tar zxvf apache-couchdb-1.6.1.tar.gz
```

Configuring CouchDB

To configure CouchDB, do the following:

- Browse to the home folder of CouchDB .
- Login as superuser.
- Configure using ./configure prompt as shown below:

```
$ cd apache-couchdb-1.6.1
$ su
Password:
# ./configure --with-erlang=/usr/lib64/erlang/usr/include/
```

It gives you the following output similar to that shown below with a concluding line saying - **You have configured Apache CouchDB, time to relax.**

```
# ./configure --with-erlang=/usr/lib64/erlang/usr/include/
checking for a BSD-compatible install... /usr/bin/install -c
checking whether build environment is sane... yes
checking for a thread-safe mkdir -p... /bin/mkdir -p
checking for gawk... gawk
checking whether make sets $(MAKE)... yes
checking how to create a ustar tar archive... gnutar
.....
.....
config.status: creating var/Makefile
config.status: creating config.h
config.status: config.h is unchanged
config.status: creating src/snappy/google-snappy/config.h
config.status: src/snappy/google-snappy/config.h is unchanged
config.status: executing depfiles commands
config.status: executing libtool commands

You have configured Apache CouchDB, time to relax.

Run `make && sudo make install' to install.
```

Installing CouchDB

Now type the following command to install CouchDB in your system.

15




```
# make && sudo make install
```

It installs CouchDB in your system with a concluding line saying - **You have installed Apache CouchDB, time to relax.**

Starting CouchDB

To start CouchDB, browse the directory in CouchDB home folder of CouchDB and use the following command:

```
$ cd apache-couchdb-1.6.1
$ cd etc
$ couchdb start
```

It starts CouchDB giving the following output:

```
Apache CouchDB 1.6.1 (LogLevel=info) is starting.
Apache CouchDB has started. Time to relax.
[info] [<0.31.0>] Apache CouchDB has started on http://127.0.0.1:5984/
[info] [<0.112.0>] 127.0.0.1 - - GET / 200
[info] [<0.112.0>] 127.0.0.1 - - GET /favicon.ico 200
```

Verification

Since CouchDB is a web interface, try to type the following homepage url in the browser.

```
http://127.0.0.1:5984/
```

It produces the following output:

```
{
  "couchdb":"Welcome",
  "uuid":"8f0d59acd0e179f5e9f0075fa1f5e804",
  "version":"1.6.1",
  "vendor":{"
    "name":"The Apache Software Foundation",
```

16

```
"version": "1.6.1"  
  
}  
  
}
```

End of ebook preview
If you liked what you saw...
Buy it from our store @ <https://store.tutorialspoint.com>