



TokuDB[®] v5.2.7 with Fractal Tree[®] Indexing
for
MySQL[®] v5.1.52
Quick Start Guide for linux
January 19, 2012

Notice: Tokutek and TokuDB are registered trademarks and Fractal Tree is a trademark of Tokutek, Inc. MySQL and Sun Microsystems are registered trademarks of Sun Microsystems, Inc.. InnoDB and Oracle are registered trademarks of Oracle Corporation.
Revision: 39012

1 Getting Started

This guide provides instructions for installing and getting started with TokuDB for MySQL®.

TokuDB, our flagship product, is a highly scalable, zero maintenance downtime, MySQL storage engine that delivers indexing-based query acceleration and enables hot schema modifications. TokuDB is a "drop-in" storage engine requiring no changes to MySQL applications or code and is fully ACID and MVCC compliant.

Providing near seamless compatibility for MySQL applications, data is loaded, inserted, and queried using standard MySQL commands, with no restrictions or special requirements. Tables can be individually defined to use TokuDB, MyISAM, InnoDB® or other MySQL-compliant storage engines.

Additional features unique to TokuDB include:

- Hot schema changes:
 - Hotindex creation: TokuDB tables support insertions, deletions and queries with no down time while indexes are being added to that table
 - Hot column addition and deletion: TokuDB tables support insertions, deletions and queries with minimal down time when an alter table adds or deletes columns.
- Fast recovery time (seconds, not hours or days)
- Immunity to database disorder or "aging" and the gradual performance degradation it causes
- 5x-15x data compression



2 System and Hardware Requirements

Operating Systems: TokuDB 5.2.7 is currently supported on 64-bit Linux only.

- CentOS 5.1 through CentOS 5.6 (64 bit).

Formal support is provided for CentOS. TokuDB for MySQL is expected to work on other Linux distributions, although support will be provided only on a best effort basis.

A virtual machine image is available for evaluations on Windows and Mac, please contact us at support@tokutek.com for more information.

Processor Architecture: TokuDB runs on the following architectures.

- X86_64

Please contact us at support@tokutek.com if you are interested in other architectures.

Memory: TokuDB Requires at least 1GB of main memory.

Disk space and configuration: Please make sure to allocate enough disk space for data, indexes and logs. In our users' experience, TokuDB typically achieves a 5x-15x space savings on data and indexes over InnoDB due to aggressive compression. TokuDB 5.2.7 does not compress logs, although it uses considerably less disk for logs than releases before 4.0.0. See the User's Guide section on "What's New" for more information.

3 Installation

Tokutek provides a binary version of MySQL 5.1.52 with a few minor changes to support TokuDB. TokuDB, MyISAM and InnoDB are included in the distribution files. The Tokutek distribution includes both the MySQL server and client for your system.

If you are installing the Operating System from scratch, the installer will give you the option of installing MySQL. Make sure that this option is not checked, or else you will subsequently need to uninstall the MySQL package that came with the Operating System.

The installation of MySQL from a tarball release file is described nicely in the MySQL 5.1 reference manual: <http://dev.mysql.com/doc/refman/5.1/en/installing-binary.html>. The following discussion includes information specific to TokuDB.

3.1 Upgrading from Previous Releases of TokuDB

It is good practice to back up your data before installing any new version of MySQL. You should protect your data by making a backup.

You must dump the database tables due to an on-disk format change. Instructions for dumping and reloading data from a table can be found in the MySQL 5.1 reference manual: <http://dev.mysql.com/doc/refman/5.1/en/>.

3.2 Removing Previous MySQL Instances

Moving tables from another storage engine can be done with a dump and reload, or more directly with an alter table (see Section 4.3).

Next, any running mysqld must be shut down using the command:

```
# mysqladmin shutdown
```

The “/etc/my.cnf” file contains configuration parameters for MySQL. It should be saved before any previous releases are removed and customized for the new release.

Once the mysql server is no longer running, any previous installation of MySQL should be uninstalled. The details of how this will happen on any particular system may vary, but for example, any MySQL installed with as an rpm can be uninstalled through rpm. To remove all installed packages with the string mysql in the name, try

```
# rpm --erase --nodeps `rpm -qa | grep -i mysql`
```

Once the system has no existing MySQL installations, it is safe to proceed with the TokuDB installation.

3.3 Downloading

To get started, navigate to Tokutek’s Support Page:

```
http://www.tokutek.com/support
```

login and download this file to your download directory:



```
mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64.tar.gz
```

After downloading, optionally verify the MD5 checksum by comparing the results of the following command to the MD5 checksum on the support page:

```
$md5sum mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64.tar.gz
```

3.4 Using data from an existing database

It is possible to use the data directory from an existing compatible database with TokuDB if installing to `/usr/local` or in a private directory. When possible make a backup of your database after performing the initial shutdown operation. The exact steps are as follows:

1. Shutdown your existing database.

```
# mysqladmin shutdown
```

2. Follow step 1, "Unpacking the tarball" of the Installation Instructions
3. Replace the data directory of your untar operation with your existing data directory
4. Follow step 3, "Run MySQL" of the Installation Instructions
5. Connect to MySQL

```
# /usr/local/mysql/bin/mysql
```

6. Install the required plugins

```
> install plugin tokudb soname 'ha_tokudb.so';  
> install plugin tokudb_user_data soname 'ha_tokudb.so';  
> install plugin tokudb_user_data_exact soname 'ha_tokudb.so';
```

7. Verify the installation of the storage engine as per "Verify the Installation"

3.5 Installation Instructions

Note: it is necessary to start TokuDB with the `mysqld_safe` script that we provide in our tarball as our version sets the `LD_LIBRARY_PATH` and `LD_PRELOAD` environment variables to enable loading the storage engine shared libraries.

We identify three common scenarios and give quick guides for installation in those cases.

3.5.1 How to install TokuDB in `/usr/local`

1. Unpacking the tarball

```
# cd /usr/local  
# rm mysql  
# tar xzf mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64.tar.gz  
# ln -s mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64 mysql
```

2. Initialize MySQL



```
# cd /usr/local/mysql
# bin/mysql_install_db --user=mysql --basedir=$PWD
```

3. Run MySQL

```
# /usr/local/mysql/bin/mysqld_safe --user=mysql
```

3.5.2 How to install TokuDB in a private directory

Note: You must execute the following as a user, not root. Executing the following as root can cause errors. If you see `ENOPERM` or `errno 13` error messages, then check permissions on the tarballs, the directories, and the files, and be sure that you are running as a user, not as root.

1. Unpack the tarball

```
$ tar xzf mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64.tar.gz
```

2. Initialize MySQL

```
$ cd mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64
$ bin/mysql_install_db --basedir=$PWD
```

3. Run MySQL

```
$ bin/mysqld_safe --basedir=$PWD
```

3.5.3 How to install TokuDB in a MySQL sandbox

Note: You must execute the following as a user, not root. Executing the following as root can cause errors. If you see `ENOPERM` or `errno 13` error messages, then check permissions on the tarballs, the directories, and the files, and be sure that you are running as a user, not as root.

1. Make the sandbox (in `$HOME/sandboxes`) and start the mysqld server

```
$ make_sandbox mysql-5.1.52-tokudb-5.2.7-38674-linux-x86_64.tar.gz
```

3.6 Verify the Installation

Verify the TokuDB Storage Engine Installation: Start a mysql client session in a command line window to verify the TokuDB storage engine installation.

Notes:

1. If you are using the mysql sandbox, start the mysql client by going into the sandbox directory (which will be a subdirectory of `$HOME/sandboxes` whose name starts with `msb`) and executing:

```
./use
```

2. If you are not using the mysql sandbox, then start the mysql client with the following command:

```
$ /usr/local/mysql/bin/mysql
```

(Note, if you are not using the mysql sandbox the default socket for mysqld is `/tmp/mysql.sock`.)



Now that the mysql client is running, you can execute the following mysql commands to query the server.

```
mysql> show engines;
```

You should see a line that reads:

```
| TokuDB | YES | Tokutek TokuDB Storage Engine | YES | NO | YES |
```

Verify the Version Numbers: For version information type:

```
mysql> show variables like '%version%';
```

On Linux, you should see:

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| protocol_version | 10 |
| tokudb_version | 5.2.7-38674 |
| version | 5.1.52 |
| version_comment | MySQL Community Server (GPL) (Tokutek 5.1.52-tokudb-5.2.7-38674) |
| version_compile_machine | x86_64 |
| version_compile_os | unknown-linux-gnu |
+-----+-----+
6 rows in set (0.04 sec)
```



3.7 Other Setups

Here, we have described common cases for installing MySQL with TokuDB. Please refer to the MySQL documentation for other issues, such as automatically starting and stopping MySQL, installing MySQL in other locations, etc. See: <http://dev.mysql.com/doc>

4 Creating Tables and Loading Data

4.1 Creating TokuDB Tables

TokuDB tables are created the same way as other tables in MySQL by specifying `engine=TokuDB` in the table definition. For example,

```
create table TABLENAME (
    ID int(11) not null
) engine=TokuDB;
```

creates a table with name TABLENAME with a single column ID and uses the TokuDB storage engine to store its data.

4.2 Loading Data

Once TokuDB tables have been created, data can be inserted or loaded using standard MySQL insert or bulk load operations. For example:

```
load data infile DATA_FILE into table TABLE_NAME;
```

loads the data from DATA_FILE into the table TABLE_NAME. See the MySQL 5.1 reference manual for more information on loading data. Please see the user guide for details on how to manage the log sizes in this case. Please see the TokuDB Users Guide for details on how to manage the log size in this case.

4.3 Migrating Data from an Existing Database

To convert an existing table from another storage engine to use the TokuDB engine,

```
alter table <tablename> engine=TokuDB;
```

If you prefer to retain the original table or to manage log size more effectively, you can dump and re-load data into a TokuDB table instead. See the FAQ section in the TokuDB Users Guide for details on dumping and re-loading and for details on managing the log size.

5 Additional Questions and Support

Tokutek looks forward to your feedback on this product and we encourage you to contact us by e-mail at support@tokutek.com. Our Technical Services team monitors this email address to ensure prompt and personal service.



The TokuDB User Guide and Frequently Asked Questions (FAQ) are available on our website at <http://www.tokutek.com/resources/product-docs>.

Our Evaluation Guide will help you understand all the capabilities of TokuDB and how they apply to your specific use-case, it can be downloaded at <http://www.tokutek.com/resources/tokudb-eval-guide>.

Visit us at <http://www.tokutek.com/support> for more information.