



How To Migrate from MyISAM



Migrating from MySQL/MyISAM to Infobright

Many users of MySQL turn to Infobright as their data volumes and analytic needs grow since Infobright offers exceptional query performance for analytic applications against large amounts of data. Migrating between MySQL's MyISAM storage engine, or other MySQL storage engines, to Infobright's column-oriented database is quite straightforward.

Infobright contains a bundled version of MySQL and installing Infobright installs a new instance of MySQL along with Infobright's Optimizer, Knowledge Grid, the Infobright Loader and the underlying columnar storage architecture. This installation also includes MySQL's MyISAM storage engine.

Unlike other storage engines that work with MySQL, it is not necessary to have an existing MySQL installation nor can Infobright be added to an existing MySQL Server installation. When installing Infobright, the assumption is that any previously existing MySQL or MyISAM database will exist in a separate installation of MySQL, installed in a different directory with a unique data path, configuration files, socket and port values. Infobright has created a Unix-based shell script to assist with this scenario. Installation guides are available at <http://www.infobright.org/wiki>.

Specify the Database

As with all MySQL databases, Infobright has a unique storage type, or engine type that must be specified when creating tables within the Infobright schema by adding the following: `ENGINE=BRIGHTHOUSE`.

Schema

Infobright neither needs nor allows the manual creation of performance structures with duplicated data such as indexes or table partitioning based on expected usage patterns of the data. When preparing the MySQL schema definition for execution in Infobright, the first thing to do is simplify the schema. This means removing all references to indexes and other constraints expressed as indexes including `PRIMARY` and `FOREIGN KEYS`, and `UNIQUE` and `CHECK` constraints.

In addition, due to Infobright's extremely high query performance levels on large volumes of data, one should consider removing all aggregate, reporting and summary tables that may be in the data model as they are unnecessary.

Data Types

Infobright supports a large subset of MySQL data types and DDL. However in some cases, inconsistencies exist between Infobright and MyISAM tables in data type support, e.g. UNSIGNED INTEGERS and the AUTO_INCREMENT attribute. These and other cases are being addressed in development for upcoming Infobright releases. Detailed information about supported data types is available at

http://www.infobright.org/wiki/Supported_Data_Types_and_Values/

Data Loading

Finally, the last step is the physical transportation of the data. Infobright provides migration utilities for this purpose, and several other migration tools exist with which Infobright is compatible. ETL tools can also satisfy the task of one-time data migrations and Infobright has created high-speed connectors for these popular ETL tools: 1. Pentaho Data Integration (PDI), and 2. JasperETL from JasperSoft (which is OEMed from Talend) or Talend Open Studio.

But for a one-time data migration, the simplest method might be the best, which is exporting the data from MyISAM tables and loading into Infobright with the Infobright Loader. Both "data out" and "data in" methods use standard MySQL syntax.

Detailed information about data loading syntax is available at

http://www.infobright.org/wiki/Data_Loading/

Database Maintenance

Once loaded an administrator can interact with the database just as they would any other in MySQL. Commands such as SHOW DATABASES, USE <database>, SHOW TABLES, SHOW TABLE STATUS, SELECT COUNT(*) FROM TABLE, etc., all work as expected.