

## ICE Kettle Connector

### Infobright Community Edition (ICE) Kettle Connector

March 11, 2009

The Infobright Kettle connector is a Kettle transformation step that is designed to load data into an Infobright table.

The connector is supported on Linux. ICE is currently also supported on Windows, and support for the connector on Windows will be added in future.

### Prerequisites

Install Kettle (Pentaho Data Integration 3.x). The connector is not compatible with Kettle 2.x.

- Install the MySQL JDBC driver, release 5, into libext/JDBC directory under your Kettle installation. The driver can be downloaded from <http://dev.mysql.com/downloads/connector/j/5.1.html>
- Delete the old mysql-connector-java-3.1.14-bin.jar as it is not compatible with Infobright

### Installation

Unpack the infobright\_kettle.zip into your Kettle directory.

### Usage

The connector is found in the components selection under **Output->BrightHouse loader**.



Figure 1: Connector dialog

The connector dialog accepts the following parameters:

## ICE Kettle Connector

- Step name
- Server host name (must be localhost currently).
- Database name
- Port number
- Username
- Password
- Infobright product
- Target schema (leave blank if same as Database name)
- Target table: The table name to load

Select the Infobright product version that corresponds to the version you are using. The options are: Community (ICE), or Enterprise (IEE).

The table must be an Infobright table (ENGINE=BRIGHTHOUSE).

**Note:** The connector works only on the local machine (host = localhost) because it uses LOAD DATA INFILE instead of LOAD DATA LOCAL INFILE. It needs to find the pipe to load on the local filesystem. Remote loading will be supported in a future release.

The list of columns passed to the connector must exactly match the column names and data types in the underlying table, and in the same order. Use a Select Values step in front of the connector to achieve this.

Under Linux, the FIFO is created in /tmp and has permissions 644. This is to ensure that mysqld can read the pipe. If this is a security concern then run Kettle under the same userid as mysqld (normally mysql).

### Known Issues

Under very rare circumstances, if the transformation is killed, it may leave the LOAD DATA INFILE process running. This will leave the table locked, so queries will hang up. This situation will only occur if the transformation is killed when the load is being initiated. However, since this only happens at the beginning of the load, the probability of this problem occurring is unlikely.

In the rare case where the transformation process has been killed but the LOAD DATA INFILE process is still running, the issue can be resolved by:

1. Going into the MySQL console and do a "show processlist;"
2. Doing a "kill <session#>;" where session# is the session of the LOAD DATA INFILE process.