

ALTIBASE® HDB™ Tools & Utilities

# Altibase Hadoop Connector User's Manual

Release 6

May 28, 2015



---

ALTIBASE HDB Tool & Utilities Altibase Hadoop Connector User's Manual

Release 6

Copyright © 2001~2015 Altibase Corporation. All rights reserved.

This manual contains proprietary information of Altibase Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright patent and other intellectual property law. Reverse engineering of the software is prohibited.

All trademarks, registered or otherwise, are the property of their respective owners.

Altibase Corporation

10F, Daerung PostTower II, 182-13,

Guro-dong Guro-gu Seoul, 152-847, South Korea

Telephone: +82-2-2082-1000 Fax: 82-2-2082-1099

Homepage: <http://www.altibase.com>

---

# Contents

<b>Preface .....</b>	<b>v</b>
About This Manual .....	vi
Audience.....	vi
Software Environment.....	vi
How This Manual is Structured.....	vi
Documentation Conventions .....	vii
Related Reading .....	viii
Online Manuals .....	viii
Altibase Welcomes Your Comments .....	viii
<b>1. Introduction to Altibase Hadoop Connector .....</b>	<b>1</b>
1.1 Background Knowledge.....	2
1.1.1 Hadoop .....	2
1.1.2 Sqoop .....	2
1.2 What is the Altibase Hadoop Connector? .....	3
<b>2. Installing the Altibase Hadoop Connector.....</b>	<b>5</b>
2.1 Software Requirements.....	6
2.1.1 Installing Sqoop.....	6
2.1.2 Installing the JDBC Driver .....	6
2.2 Installing the Altibase Hadoop Connector .....	7
2.3 Executing & Testing .....	8
<b>3. Functions .....</b>	<b>9</b>
3.1 Command-line Options.....	10
3.1.1 Syntax.....	10
3.1.2 Options .....	10
3.2 Import .....	11
3.2.1 Importing to HDFS in a Text File.....	11
3.2.2 Importing to HDFS in a Sequence File .....	12
3.2.3 Importing to HDFS in an Avro File .....	12
3.2.4 Importing to HDFS Using a Query .....	12
3.2.5 Importing to Hive.....	12
3.3 Export.....	14
3.3.1 Inserting Data .....	14
3.3.2 Inserting Data From a CSV File .....	14
3.3.3 Inserting Data in Batch Mode .....	14
3.3.4 Updating Data.....	15
3.3.5 Updating or Inserting Data.....	15
3.3.6 Inserting Data Using a Staging-Table .....	15
3.4 list-databases .....	16
3.5 list-tables.....	17



# Preface

---

# About This Manual

This manual explains how to use the Altibase Hadoop Connector with ALTIBASE HDB/XDB and Hadoop.

## Audience

This manual has been prepared for the following ALTIBASE HDB users:

- Database Administrators
- Database Users
- Application Program Developers
- Technical Assistance Team

This manual assumes that you have the following background :

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides.
- Some experience working with relational databases or exposure to database concepts
- Some experience with computer programming
- Some experience with database server administration, operating system administration or network administration

## Software Environment

This manual was written under the assumption that ALTIBASE HDB version 6.0 or above, or ALTIBASE XDB version 6.0 or above is being used as a database server.

## How This Manual is Structured

This manual covers the following topics :

- [Chapter1: Introduction to Altibase Hadoop Connector](#)  
This chapter introduces the Altibase Hadoop Connector and provides background knowledge.
- [Chapter2: Installing the Altibase Hadoop Connector](#)  
This chapter explains how to install the Altibase Hadoop Connector and required software.
- [Chapter3: Functions](#)  
Altibase Hadoop Connector functions are explained.
- [Appendix A. Data Type](#)

This appendix describes the data types supported by the Altibase Hadoop Connector.

## Documentation Conventions

This section offers documentation conventions as follows. They make it easier to gather information from Altibase manuals.

- Sample Code Conventions

### Sample Code Conventions

The code examples explain SQL statements, stored procedures, iSQL statements, and other command line syntax. The following table describes the printing conventions used in the code examples.

Convention	Meaning	Example
[ ]	Indicates an optional item.	VARCHAR [(size)] [[FIXED  ] VARIABLE]
{ }	Indicates a mandatory field for which one or more items must be selected.	{ ENABLE   DISABLE   COMPILE }
	A delimiter between optional or mandatory arguments.	{ ENABLE   DISABLE   COMPILE } [ ENABLE   DISABLE   COMPILE ]
. . . .	Indicates that the previous argument is repeated, or that sample code has been omitted.	iSQL> select e_lastname from employees; E_LASTNAME ----- Moon Davenport 
Other symbols	Symbols other than those shown above are part of the actual code.	EXEC :p1 := 1; acc NUMBER(11,2);
Italics	Statement elements in italics indicate variables and special values specified by the user.	SELECT * FROM table_name; CONNECT <i>userID/password</i> ;
Lower Case Letters	Indicate program elements set by the user, such as table names, column names, file names, etc.	SELECT e_lastname FROM employees;
Upper Case Letters	Keywords and all elements provided by the system appear in upper case.	DESC SYSTEM_.SYS_INDICES_;

## Related Reading

For additional technical information, consult the following manuals.

- Altibase Installation Guide
- Altibase Getting Started
- Altibase Administrator's Manual
- Altibase General Reference
- Altibase Error Message Reference

## Online Manuals

Online versions of our manuals (PDF or HTML) are available from Altibase's Customer Support site(<http://support.altibase.com/>).

## Altibase Welcomes Your Comments

Please let us know what you like or dislike about our manuals. To help us with future versions of our manuals, please tell us about any corrections or classifications that you would find useful.

Include the following information :

- The name and version of the manual that you are using
- Any comments that you have about the manual
- Your name, address, and phone number

When you need an immediate assistance regarding technical issues, please contact Altibase's Customer Support site(<http://support.altibase.com/>).

Thank you. We appreciate your feedback and suggestions.



# 1 Introduction to Altibase Hadoop Connector

---

This chapter introduces the Altibase Hadoop Connector and provides background knowledge.

## 1.1 Background Knowledge

This chapter introduces the basic concept of transferring data between Hadoop and ALTIBASE HDB/XDB with Sqoop.

### 1.1.1 Hadoop

Hadoop is a system that is suitable for the analysis and management of large amounts of data. Hadoop is one of the solutions seeing increased interest as demands for parallel processing of data warehouses and the management of 'big data' increase and the popularity of cloud and distributed computing grows.

Hadoop is open source software based on Java, and is a framework which distributively processes large amounts of data spread across clusters composed of multiple computers.

Hadoop is composed of HDFS(Hadoop Distributed File System) and MapReduce.

### 1.1.2 Sqoop

Sqoop is a tool for data transfer between Hadoop and relational databases, and is open source software. Using Sqoop, the user can import data from RDBMS to HDFS and export it back to RDBMS.

For the import and export of data, Sqoop uses MapReduce. MapReduce offers Fault Tolerance and Parallel Operation.

## 1.2 What is the Altibase Hadoop Connector?

The Altibase Hadoop Connector enables the efficient transfer of data between Hadoop and ALTIBASE HDB/XDB servers. It also allows ALTIBASE HDB/XDB to handle data management and Hadoop to handle data analysis.

In other words, the Altibase Hadoop Connector lets the user connect to the HDB server and export data into HDFS or HIVE for the purpose of processing data from Hadoop.

The Altibase Hadoop Connector operates on Sqoop and supports nearly all functions provided by it. Also, users that have experience with Sqoop will easily be able to use the Altibase Hadoop Connector, as it uses a command line argument structure similar to Sqoop.

## 1.2 What is the Altibase Hadoop Connector?

# **2** Installing the Altibase Hadoop Connector

---

This chapter explains how to install the Altibase Hadoop Connector and required software.

## 2.1 Software Requirements

The following software must be installed in order to install and run the Altibase Hadoop Connector.

- JRE (Java Runtime Environment) or JDK (Java Development Kit) version 1.6 or higher
- Sqoop version 1.4.4 or higher
- ALTIBASE HDB version 5.0 or higher, ALTIBASE XDB version 6.0 or higher

### 2.1.1 Installing Sqoop

The Altibase Hadoop Connector is based on Apache Sqoop, which is distributed under the Apache License version 2.0.

Install Sqoop in the following order and install the Altibase JDBC driver into the Sqoop environment to enable sqoop to connect to ALTIBASE HDB/XDB.

1. Download the Sqoop (package Sqoop-1.4.4.bin\_hadoop-1.0.0.tar.gz) that supports Hadoop 1.0 from <http://Mirror.apache-kr.org/sqoop/1.4.4>. The current Altibase Hadoop Connector officially supports Sqoop 1.4.4 (Sqoop 1.9 will be supported soon).
2. Install Sqoop using the installation manual from the Sqoop homepage.

### 2.1.2 Installing the JDBC Driver

Install the JDBC driver as follows after installing Sqoop.

#### 2.1.2.1 When using ALTIBASE HDB

Copy the JDBC driver file for ALTIBASE HDB (`$ALTIBASE_HOME/lib/Altibase.jar`) to the `$SQOOP_HOME/lib`.

```
% cp $ALTIBASE_HOME/lib/Altibase.jar $SQOOP_HOME/lib
```

#### 2.1.2.2 When using ALTIBASE XDB

Copy the JDBC driver file for ALTIBASE XDB (`$ALTIBASE_XDB_HOME/lib/Altibase_xdb.jar`) to the `$SQOOP_HOME/lib`.

```
% cp $ALTIBASE_XDB_HOME/lib/Altibase_xdb.jar $SQOOP_HOME/lib
```

## 2.2 Installing the Altibase Hadoop Connector

Install the Altibase Hadoop Connector in the following order:

1. Download the Altibase Hadoop Connector Package from <http://support.altibase.com/>.

Package for HDB: `altibase_sqoop14_connector.jar`

Package for XDB: `altibase_xdb_sqoop14_connector.jar`

2. Copy the Altibase Hadoop connector package to `$SQOOP_HOME/lib`.

### 2.3 Executing & Testing

The following is a command which retrieves the table list from the ALTIBASE HDB server and can check whether or not the Altibase Hadoop Connector is installed properly.

```
% sqoop list-tables
--connect jdbc:Altibase://127.0.0.1:20300/mydb
--driver Altibase.jdbc.driver.AltibaseDriver
--username SYS
--password MANAGER
--connection-manager com.altibase.sqoop.manager.AltibaseManager
```

If the following table list is printed out with the log below after execution, it means that the Altibase Hadoop Connector has been loaded properly.

```
13/10/02 13:48:15 INFO manager.AltibaseManager: init default option autocommit false
13/10/02 13:48:15 INFO manager.SqlManager: Using default fetchSize of 1000
13/10/02 13:48:15 INFO manager.AltibaseManager: Altibase HDB manager 1.0 connector create
```



# 3 Functions

---

This chapter explains Altibase Hadoop Connector functions with examples.

## 3.1 Command-line Options

All Sqoop features and functions can be found at:

<http://sqoop.apache.org/docs/1.4.4/SqoopUserGuide.html>

### 3.1.1 Syntax

Listed below is an example of the Altibase Hadoop Connector syntax:

```
sqoop <command> --connect <url>
--driver <driver>
--username <user>
--password <password>
--connection-manager com.altibase.sqoop.manager.AltibaseManager
```

### 3.1.2 Options

Options	Explanation
<command>	Specifies the command: export import list-databases list-tables
--connect <url>	Specifies the ALTIBASE HDB or XDB URL: HDB: jdbc:Altibase://ip:port/mydb XDB: jdbc:Altibase:remote://ip:port/mydb
--driver <driver>	Specifies the ALTIBASE HDB or XDB JDBC driver class: HDB: Altibase.jdbc.driver.AltibaseDriver XDB: Altibase_xdb.jdbc.driver.AltibaseDriver
--username <user>	Specifies the user
--password <password>	Specifies the password.
--connection-manager <connector>	Specifies the Altibase Hadoop Connector: HDB: com.altibase.sqoop.manager.AltibaseManager XDB: com.altibase_xdb.sqoop.manager.Altibasemanager

Additional commands, import and export options can be found at: <http://sqoop.apache.org/docs/1.4.4/SqoopUserGuide.html>.

## 3.2 Import

Import is a function which imports the data of ALTIBASE HDB or XDB to HDFS or Hive. This section provides explanations of these functions with examples.

### 3.2.1 Importing to HDFS in a Text File

The following options can be used with the Sqoop command to store data of a certain table existing in ALTIBASE HDB to a configured directory on HDFS in text file format.

```
% sqoop import
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  -split-by <split_name>
```

By default, text files are stored as CSV. Separators, such as field termination characters, line termination characters can be specified with the options listed in the following table. If the user does not specify one of the options, default values are taken.

Option	Description	Default Value
--enclosed-by	The character which encloses the field; this option is required.	\"
--escaped-by	The escape character	
--fields-terminated-by	The field separator character	,
--lines-terminated-by	The line termination character	\n
--mysql-delimiters	The MySQL default delimiter; the following are available: fields-terminated-by: , lines-terminated-by: \n escaped-by: \ optionally-enclosed-by:	
--optionally-enclosed-by	The character which encloses the field; this option is optional.	

*Note: If a line separator or a field separator is included in the field, it must be enclosed in punctuation marks, such as quotation marks(""). If a quotation mark is included in the field, an escape character must be prefixed to the quotation mark to indicate that it is not a punctuation mark. Due to a sqoop bug found at the time of writing(sqoop operates abnormally when the punctuation mark and the escape character are duplicate), the Altibase Hadoop Connector only provides quotation marks as the default value for punctuation marks and offers support for the user to directly specify the escape character.*

### 3.2.2 Importing to HDFS in a Sequence File

The following options can be used with the Sqoop command to store data of a certain table existing in ALTIBASE HDB to a specified directory on HDFS in sequence file format.

```
% sqoop import
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --split-by <split_name>
  --as-sequencefile
```

### 3.2.3 Importing to HDFS in an Avro File

The following options can be used with the Sqoop command to store data of a certain table existing in ALTIBASE HDB to a specified directory on HDFS in Avro file format.

```
% sqoop import
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --split-by <split_name>
  --as-avrodatafile
```

### 3.2.4 Importing to HDFS Using a Query

The following options can be used with the Sqoop command to store data of a certain table existing in ALTIBASE HDB into HDFS using a user-specified query statement.

```
% sqoop import
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --split-by <split_name>
  --boundary-query <query>
```

### 3.2.5 Importing to Hive

The following options can be used with the Sqoop command to store data of a certain table existing in ALTIBASE HDB into Hive.

```
% sqoop import
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
```

```
--connection-manager com.altibase.sqoop.manager.AltibaseManager  
-table <table_name>  
-split-by <split_name>  
--hive-import
```

# 3.3 Export

Export is a function which exports HDFS data to ALTIBASE HDB or XDB. Depending on the option, data is inserted into the database or existing data is updated.

## 3.3.1 Inserting Data

The following options can be used with the Sqoop command to store HDFS data in a certain table existing in ALTIBASE HDB.

```
% sqoop export
  -D sqoop.export.records.per.statement=<size>
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --export-dir <dir>
```

If the value set for `sqoop.export.records.per.statement` is larger than 1, the Altibase Hadoop Connector operates in batch mode. Batch mode inserts multiple records in one EXECUTE statement. This option must precede other options.

If the value set for `sqoop.export.records.per.statement` is 1, it is possible that the Altibase Hadoop Connector will not operate in batch mode.

## 3.3.2 Inserting Data From a CSV File

When exporting with a CSV file, options which specify delimiters are available for use. For further information on options available for use, please refer to "Importing to HDFS in a Text File".

## 3.3.3 Inserting Data in Batch Mode

The following options can be used with the Sqoop command to store HDFS data to a certain table existing in ALTIBASE HDB using the batch mode feature of the Altibase JDBC driver.

```
% sqoop export
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --export-dir <dir>
  --batch
```

Regardless of whether or not the batch option has been specified, the Altibase Hadoop Connector runs in batch mode and 100 records are inserted by default for one Execute operation. To disable the Altibase Hadoop Connector from running in batch mode, the `-D sqoop.export.records.per.statement` option should be set to 1.

### 3.3.4 Updating Data

The following options can be used with the Sqoop command to update data of a certain table existing in ALTIBASE HDB to HDFS data.

```
% sqoop export
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --export-dir <dir>
  -update-key <column>
```

### 3.3.5 Updating or Inserting Data

This function performs an update if HDFS data exists in ALTIBASE HDB, and performs an insert if not. The following options can be used with the Sqoop command:

```
% sqoop export
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --export-dir <dir>
  -update-key <column>
  --update-mode allowinsert
```

*Caution: Since this function uses the MERGE statement of ALTIBASE HDB, it is only supported when the Altibase Hadoop Connector is used with ALTIBASE HDB version 6.3.1 or higher.*

### 3.3.6 Inserting Data Using a Staging-Table

Since Sqoop exports with multiple transactions, some data can fail to be committed. The --staging-table option is used to prevent such failure and first inserts HDFS data to the table specified for this option, and then moves this data to a target table in ALTIBASE HDB. The following options can be used with the Sqoop command:

```
% sqoop export
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
  -table <table_name>
  --export-dir <dir>
  -update-key <column>
  --staging-table <table_name>
```

## 3.4 list-databases

The following options can be used with the Sqoop command to retrieve the database list of ALTI-BASE HDB.

```
% sqoop list-databases
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
```



## 3.5 list-tables

The following options can be used with the Sqoop command to retrieve the tables existing in ALTI-BASE HDB.

```
% sqoop list-tables
  --connect <url>
  --driver <jdbc_driver>
  --username <user>
  --password <password>
  --connection-manager com.altibase.sqoop.manager.AltibaseManager
```

3.5 list-tables

# Appendix A. Data Type

---

This appendix describes data types supported by Altibase Hadoop Connector.

## Supported Data Type

The following is the table that shows the conversion of data type between Sqoop and ALTIBASE HDB/XDB when importing and exporting using the Altibase Hadoop Connector. Also it shows whether it supports for import or export depending on the data types

HDB Data Type	XDB Data Type	Sqoop Data Type	Import	Export
CHAR	CHAR	String	O	O
VARCHAR	VARCHAR	String	O	O
NCHAR	NCHAR	String	O	O
NVARCHAR	NVARCHAR	String	O	O
INTEGER	INTEGER	Integer	O	O
BIGINT	BIGINT	Long	O	O
SMALLINT	SMALLINT	Integer	O	O
NUMBER	NUMBER	Double	O	O
NUMERIC	NUMERIC	java.math.BigDecimal	O	O
DECIMAL	DECIMAL	java.math.BigDecimal	O	O
FLOAT	FLOAT	Double	O	O
DOUBLE	DOUBLE	Double	O	O
REAL	REAL	FLOAT	O	O
DATE	DATE	java.sql.Timestamp	O	O
BLOB	BLOB	com.cloudera.sqoop.lib.BlobRef	O	X
CLOB	CLOB	com.cloudera.sqoop.lib.ClobRef	O	X

*Caution: BLOB and CLOB data types are supported only for import in Sqoop so Altibase Hadoop Connector also supports import as well. The export for BLOB and CLOB data types is to be supported soon.*

Supported Data Type

# Index

## **A**

Altibase Hadoop Connector 3  
Altibase Hadoop Connector syntax 10  
Avro File 12

## **B**

Batch Mode 14

## **C**

Command-line Options 10  
CSV File 14

## **D**

Data Type 19

## **E**

Executing 8  
Export 14

## **H**

Hadoop 2  
Hive 12

## **I**

Import 11  
Inserting Data 14  
Installing Sqoop 6  
Installing the Altibase Hadoop Connector 7  
Installing the JDBC Driver 6

## **L**

list-databases 16  
list-tables 17

## **Q**

Query 12

## **S**

Sequence File 12  
Software Requirments 6  
Sqoop 2  
Sqoop Data Type 19  
Staging-Table 15

## **T**

Testing 8  
Text File 11

## **U**

Updating Data 15  
Updating or Inserting Data 15