

Oracle® Big Data Connectors

User's Guide

Release 1 (1.1)

E36049-03

October 2012

Describes installation and use of the Oracle Big Data Connectors: Oracle Direct Connector for Hadoop Distributed File System, Oracle Loader for Hadoop, Oracle Data Integrator Application Adapter for Hadoop, and Oracle R Connector for Hadoop.

Oracle Big Data Connectors User's Guide, Release 1 (1.1)

E36049-03

Copyright © 2011, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Cloudera, Cloudera CDH, and Cloudera Manager are registered and unregistered trademarks of Cloudera, Inc.

Contents

Preface	vii
Audience	vii
Documentation Accessibility	vii
Related Documents	vii
Conventions	viii
What's New in Oracle Big Data Connectors?	ix
Changes in Release 1.1	ix
1 Getting Started with Oracle Big Data Connectors	
About Oracle Big Data Connectors	1-1
Big Data Concepts and Technologies	1-2
What is MapReduce?	1-2
What is Apache Hadoop?	1-2
Downloading the Oracle Big Data Connectors Software	1-3
Oracle Direct Connector for Hadoop Distributed File System Setup	1-3
Required Software	1-3
Installing and Configuring Hadoop	1-4
Installing Oracle Direct Connector for HDFS	1-4
Granting User Access to Oracle Direct Connector for HDFS	1-5
Oracle Loader for Hadoop Setup	1-6
Required Software	1-6
Installing Oracle Loader for Hadoop	1-6
Oracle Data Integrator Application Adapter for Hadoop Setup	1-6
System Requirements and Certifications	1-7
Technology-Specific Requirements	1-7
Location of ODI Application Adapter for Hadoop	1-7
Setting Up the Topology	1-7
Oracle R Connector for Hadoop Setup	1-7
Installing the Software on Hadoop	1-7
Providing Remote Client Access to R Users	1-9
2 Oracle Direct Connector for Hadoop Distributed File System	
About Oracle Direct Connector for HDFS	2-1
Creating an External Table for HDFS	2-2

Basic SQL Syntax for the External Table.....	2-2
Testing the External Table.....	2-3
External Table Example.....	2-3
Publishing the HDFS Data Paths	2-4
ExternalTable Command	2-4
How to Publish Data Pump Files.....	2-5
Creating a Configuration File.....	2-5
Configuration Properties	2-6
Querying Data in HDFS	2-9

3 Oracle Loader for Hadoop

What Is Oracle Loader for Hadoop?	3-1
Using Oracle Loader for Hadoop	3-2
Implementing InputFormat.....	3-2
Creating the loaderMap Document.....	3-4
Accessing Table Metadata	3-5
Invoking OraLoader	3-6
Loading Files Into an Oracle Database (Offline Loads Only).....	3-7
Output Modes During OraLoader Invocation	3-7
JDBC Output.....	3-7
Oracle OCI Direct Path Output.....	3-8
Delimited Text Output	3-9
Oracle Data Pump Output.....	3-10
Balancing Loads When Loading Data into Partitioned Tables	3-11
Using the Sampling Feature	3-11
Tuning Load Balancing and Sampling Behavior.....	3-12
Does Oracle Loader for Hadoop Always Use the Sampler's Partitioning Scheme?	3-13
What Happens When a Sampling Feature Property Has an Invalid Value?	3-13
Primary Configuration Properties for the Load Balancing Feature	3-13
OraLoader Configuration Properties	3-14
Example of Using Oracle Loader for Hadoop	3-20
Target Table Characteristics	3-21
Supported Data Types.....	3-21
Supported Partitioning Strategies.....	3-22
Loader Map XML Schema Definition	3-22
OraLoader for Hadoop Configuration Properties	3-23
Third-Party Licenses for Bundled Software	3-34
Apache Licensed Code	3-34
Apache Avro avro-1.5.4.jar	3-37
Apache Commons Mathematics Library 2.2.....	3-37
Jackson JSON Library 1.5.2.....	3-38

4 Oracle Data Integrator Application Adapter for Hadoop

Introduction	4-1
Concepts	4-1
Knowledge Modules.....	4-2
Security	4-2

Setting Up the Topology	4-3
Setting Up File Data Sources	4-3
Setting Up Hive Data Sources	4-3
Setting Up the Oracle Data Integrator Agent to Execute Hadoop Jobs	4-4
Configuring Oracle Data Integrator Studio for Executing Hadoop Jobs on the Local Agent	4-6
Setting Up an Integration Project	4-6
Creating an Oracle Data Integrator Model from a Reverse-Engineered Hive Model	4-6
Creating a Model.....	4-6
Reverse Engineering Hive Tables.....	4-6
Designing the Interface	4-7
Loading Data from Files into Hive	4-7
Validating and Transforming Data Within Hive	4-8
Loading Data into an Oracle Database from Hive and HDFS.....	4-10

5 Oracle R Connector for Hadoop

About Oracle R Connector for Hadoop	5-1
Oracle R Connector for Hadoop APIs.....	5-1
Access to Oracle Database	5-2
Scenarios for Using Oracle R Packages	5-2
Security Notes for Oracle R Connector for Hadoop	5-2
Functions in Alphabetical Order	5-3
Functions by Category	5-3
Making Connections.....	5-4
Copying Data	5-4
Exploring Files	5-4
Writing MapReduce Functions	5-4
Debugging Scripts.....	5-5
Executing Scripts	5-5
ORCH mapred.config Class	5-5
Example R Programs Using Oracle R Connector for Hadoop	5-6
Using the Examples	5-6
Using the Demos	5-10
hadoop.exec	5-13
hadoop.run	5-16
hdfs.attach	5-18
hdfs.cd	5-20
hdfs.cp	5-21
hdfs.describe	5-22
hdfs.download	5-23
hdfs.exists	5-24
hdfs.get	5-25
hdfs.id	5-26
hdfs.ls	5-27
hdfs.mkdir	5-28
hdfs.mv	5-29
hdfs.parts	5-30
hdfs.pull	5-31

hdfs.push.....	5-32
hdfs.put	5-34
hdfs.pwd	5-36
hdfs.rm.....	5-37
hdfs.rmdir	5-38
hdfs.root	5-39
hdfs.sample.....	5-40
hdfs.setroot	5-41
hdfs.size.....	5-42
hdfs.upload.....	5-43
is.hdfs.id.....	5-45
orch.connect.....	5-46
orch.dbcon.....	5-48
orch.dbg.off	5-50
orch.dbg.on	5-51
orch.dbg.output.....	5-52
orch.dbinfo	5-53
orch.disconnect	5-54
orch.dryrun	5-56
orch.export	5-57
orch.keyval.....	5-58
orch.keyvals.....	5-59
orch.pack	5-61
orch.reconnect	5-62
orch.unpack	5-63
orch.version	5-64

Index

Preface

The *Oracle Big Data Connectors User's Guide* describes how to install and use Oracle Big Data Connectors:

- Oracle Loader for Hadoop
- Oracle R Connector for Hadoop
- Oracle Direct Connector for Hadoop Distributed File System
- ODI Application Adapter for Hadoop

Audience

This document is intended for users of Oracle Big Data Connectors, including the following:

- Application developers
- Java programmers
- System administrators
- Database administrators

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents:

- *Oracle Loader for Hadoop Java API Reference*
- *Oracle Fusion Middleware Application Adapters Guide for Oracle Data Integrator*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

What's New in Oracle Big Data Connectors?

This chapter briefly describes the new features in this release of Oracle Big Data Connectors and provides links to more information.

Changes in Release 1.1

Release 1.1 of Oracle Big Data Connectors includes many new and modified functions in Oracle R Connector for Hadoop.

Name Changes

The R package name changed from `ORHC` to `ORCH`. All functions with names beginning with `orhc` are renamed to begin with `orch`.

The `orhc.which` function is now named `orch.dbcon`.

New Functions

The following functions are new in Release 1.1:

- `hdfs.describe`
- `hdfs.id`
- `hdfs.mv`
- `hdfs.root`
- `hdfs.setroot`
- `is.hdfs.id`
- `orch.dbcon`
- `orch.dbg.on`
- `orch.dbg.output`
- `orch.dbinfo`
- `orch.dryrun`
- `orch.export`
- `orch.keyval`
- `orch.keyvals`
- `orch.pack`
- `orch.unpack`
- `orch.version`

Modified Functions

The following functions have modified syntax in release 1.1:

- `hadoop.exec`
- `hadoop.run`
- `hdfs.attach`
- `hdfs.put`
- `hdfs.rm`
- `hdfs.rmdir`
- `hdfs.size`
- `orch.disconnect`

Desupported Functions

These functions are desupported in Release 1.1:

- `orch.which`: Use `orch.dbcon` to obtain a connection object.

Deprecated Functions

These functions are deprecated in release 1.1 and will be desupported in release 2.0:

- `keyval`: use `orch.keyval` to generate key-value pairs.
- `orch.reconnect`: Use `orch.connect` to reconnect using a connection object returned by `orch.dbcon`.

Getting Started with Oracle Big Data Connectors

This chapter describes the Oracle Big Data Connectors, provides installation instructions, and identifies the permissions needed for users to access the connectors.

This chapter contains the following sections:

- [About Oracle Big Data Connectors](#)
- [Big Data Concepts and Technologies](#)
- [Downloading the Oracle Big Data Connectors Software](#)
- [Oracle Direct Connector for Hadoop Distributed File System Setup](#)
- [Oracle Loader for Hadoop Setup](#)
- [Oracle Data Integrator Application Adapter for Hadoop Setup](#)
- [Oracle R Connector for Hadoop Setup](#)

About Oracle Big Data Connectors

Oracle Big Data Connectors facilitate data access between data stored in a Hadoop cluster and Oracle Database. They can be licensed for use on either Oracle Big Data Appliance or a Hadoop cluster running on commodity hardware.

These are the connectors:

- **Oracle Direct Connector for Hadoop Distributed File System:** Enables Oracle Database to access data stored in Hadoop Distributed File System (HDFS). The data can remain in HDFS, or it can be loaded into an Oracle database.
- **Oracle Loader for Hadoop:** Provides an efficient and high-performance loader for fast movement of data from a Hadoop cluster into a table in an Oracle database. Oracle Loader for Hadoop repartitions the data if necessary and transforms it into a database-ready format. It optionally sorts records by primary key before loading the data or creating output files. Oracle Loader for Hadoop is a MapReduce application that is invoked as a command-line utility. It accepts the generic command-line options that are supported by the Tool interface.
- **Oracle Data Integrator Application Adapter for Hadoop:** Extracts, transforms, and loads data from a Hadoop cluster into tables in an Oracle database, as defined using a graphical user interface.
- **Oracle R Connector for Hadoop:** Provides an interface between a local R environment, Oracle Database, and Hadoop, allowing speed-of-thought, interactive analysis on all three platforms. Oracle R Connector for Hadoop is

designed to work independently, but if the enterprise data for your analysis is also stored in Oracle Database, then the full power of this connector is achieved when it is used with Oracle R Enterprise.

Individual connectors may require that software components be installed in Oracle Database and the Hadoop cluster. Users may also need additional access privileges in Oracle Database.

See Also: My Oracle Support Information Center: Big Data Connectors (ID 1487399.2) and its related information centers.

Big Data Concepts and Technologies

Enterprises are seeing large amounts of data coming from multiple sources. Click-stream data in web logs, GPS tracking information, data from retail operations, sensor data, and multimedia streams are just a few examples of vast amounts of data that can be of tremendous value to an enterprise if analyzed. The unstructured and semi-structured information provided by raw data feeds is of little value in and of itself. The data needs to be processed to extract information of real value, which can then be stored and managed in the database. Analytics of this data along with the structured data in the database can provide new insights into the data and lead to substantial business benefits.

What is MapReduce?

MapReduce is a parallel programming model for processing data on a distributed system. It can process vast amounts of data in a timely manner and can scale linearly. It is particularly effective as a mechanism for batch processing of unstructured and semi-structured data. MapReduce abstracts lower level operations into computations over a set of keys and values.

A simplified definition of a MapReduce job is the successive alternation of two phases, the map phase and the reduce phase. Each map phase applies a transform function over each record in the input data to produce a set of records expressed as key-value pairs. The output from the map phase is input to the reduce phase. In the reduce phase, the map output records are sorted into key-value sets so that all records in a set have the same key value. A reducer function is applied to all the records in a set and a set of output records are produced as key-value pairs. The map phase is logically run in parallel over each record while the reduce phase is run in parallel over all key values.

What is Apache Hadoop?

Apache Hadoop is the software framework for the development and deployment of data processing jobs based on the MapReduce programming model. At the core, Hadoop provides a reliable shared storage and analysis system¹. Analysis is provided by MapReduce. Storage is provided by the Hadoop Distributed File System (HDFS), a shared storage system designed for MapReduce jobs.

The Hadoop ecosystem includes several other projects including Apache Avro, a data serialization system that is used by Oracle Loader for Hadoop.

Cloudera's Distribution including Apache Hadoop (CDH) is installed on Oracle Big Data Appliance. You can use Oracle Big Data Connectors on a Hadoop cluster running

¹ *Hadoop: The Definitive Guide, Third Edition* by Tom White (O'Reilly Media Inc., 2012, 978-1449311520).

CDH or the equivalent Apache Hadoop components, as described in the setup instructions in this chapter.

See Also: For conceptual information about the Hadoop technologies, the following third-party publication:

Hadoop: The Definitive Guide, Third Edition by Tom White (O'Reilly Media Inc., 2012, ISBN: 978-1449311520).

Downloading the Oracle Big Data Connectors Software

You can download Oracle Big Data Connectors from Oracle Technology Network or Oracle Software Delivery Cloud.

To download from Oracle Technology Network:

1. Use any browser to visit this website:

<http://www.oracle.com/technetwork/bdc/big-data-connectors/downloads/index.html>

2. Click the name of each connector to download a zip file containing the installation files.

To download from Oracle Software Delivery Cloud:

1. Use any browser to visit this website:

<https://edelivery.oracle.com/>

2. Accept the Terms and Restrictions to see the Media Pack Search page.

3. Select the search terms:

Select a Product Pack: Oracle Database

Platform: Linux x86-64

4. Click **Go** to display a list of product packs.
5. Select Oracle Big Data Connectors Media Pack for Linux x86-64 (B65965-0x), and then click **Continue**.
6. Click **Download** for each connector to download a zip file containing the installation files.

Oracle Direct Connector for Hadoop Distributed File System Setup

Oracle Direct Connector for Hadoop Distributed File System (HDFS) is installed and runs on the system where Oracle Database runs. Before installing Oracle Direct Connector for HDFS, verify that you have the required software.

Required Software

Oracle Direct Connector for HDFS requires the following software:

- Cloudera's Distribution including Apache Hadoop version 3 (CDH3) or Apache Hadoop 0.20.2.
- Java Development Kit (JDK) 1.6.0_8 or later for CDH3. Cloudera recommends version 1.6.0_26.
- Oracle Database release 11g release 2 (11.2.0.2 or 11.2.0.3) for Linux.

- To support the Oracle Data Pump file format, an Oracle Database one-off patch. To download this patch, go to <http://support.oracle.com> and search for bug 14557588.
- The same version of Hadoop on the Oracle Database system as your Hadoop cluster, either CDH3 or Apache Hadoop 0.20.2.
- The same version of JDK on the Oracle Database system as your Hadoop cluster.

Installing and Configuring Hadoop

Oracle Direct Connector for HDFS works as a Hadoop client. You do not need to configure Hadoop on the Oracle Database system to run MapReduce jobs for Oracle Direct Connector for HDFS. However, you must install Hadoop on the Oracle Database system and minimally configure it for Hadoop client use only.

To configure the Oracle Database system as a Hadoop client:

1. Install and configure CDH3 or Apache Hadoop 0.20.2 on the Oracle Database system. If you are using Oracle Big Data Appliance, then complete the procedures for providing remote client access in the *Oracle Big Data Appliance Software User's Guide*. Otherwise, follow the installation instructions provided by the distributor (Cloudera or Apache).
2. Ensure that Oracle Database has access to Hadoop and HDFS:
 - a. Log in to the system where Oracle Database is running by using the Oracle Database account.
 - b. Open a Bash shell and enter this command:

```
$HADOOP_HOME/bin/hadoop fs -ls /user
```

In this command, `$HADOOP_HOME` is the absolute path to the Hadoop home directory. You should see a list of files. If not, then first ensure that the Hadoop cluster is up and running. If the problem persists, then you must correct the Hadoop client configuration so that Oracle Database has access to the Hadoop cluster file system.

The Oracle Database system is now ready for use as a Hadoop client. No other Hadoop configuration steps are needed.

Installing Oracle Direct Connector for HDFS

Complete this procedure to install Oracle Direct Connector for HDFS on the Oracle Database system.

To install Oracle Direct Connector for HDFS:

1. Download the zip file to a directory on the system where Oracle Database runs.
2. Unzip `orahdfs-version.zip` into a directory. The unzipped files have the structure shown in [Example 1-1](#).
3. Open the `hdfs_stream` Bash shell script in a text editor, and make the changes indicated by the comments in the script.

The `hdfs_stream` script is the preprocessor script for the HDFS external table.

4. Run the `hdfs_stream` script from the Oracle Direct Connector for HDFS installation directory. You should see this usage information:

```
$ bin/hdfs_stream
Oracle Direct HDFS Release 1.0.0.0.0 - Production
```

Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.
 Usage: \$HADOOP_HOME/bin/hadoop jar orahdfs.jar
 oracle.hadoop.hdfs.exttab.HdfsStream <locationPath>

If not, then ensure that the operating system user that Oracle Database is running under has the following permissions:

- Read and execute permissions on the `hdfs_stream` script:

```
$ ls -l $DIRECTHDFS_HOME/bin/hdfs_stream
-rwxr-xr-x 1 oracle oinstall 2273 Apr 27 15:51 hdfs_stream
```

If you do not see these permissions, then enter a `chmod` command to fix them:

```
$ chmod 755 $DIRECTHDFS_HOME/bin/hdfs_stream
```

In these commands, `DIRECTHDFS_HOME` represents the Oracle Direct Connector for HDFS home directory.

- Read permission on `DIRECTHDFS_HOME/jlib/orahdfs.jar`.
5. Create a database directory for the `orahdfs-version/bin` directory where `hdfs_stream` resides. In this example, the Oracle Direct Connector for HDFS kit is installed in `/etc`:

```
SQL> CREATE OR REPLACE DIRECTORY hdfs_bin_path AS '/etc/orahdfs-1.0/bin'
```

Example 1–1 Structure of the orahdfs Directory

```
orahdfs-version
bin/
  hdfs_stream
jlib/
  orahdfs.jar
log/
README.txt
```

Granting User Access to Oracle Direct Connector for HDFS

Oracle Database users require these privileges to use Oracle Direct Connector for HDFS:

- CREATE SESSION
- EXECUTE on the `UTL_FILE` PL/SQL package
- READ and EXECUTE on the `HDFS_BIN_PATH` directory created during the installation of Oracle Direct Connector for HDFS. Do not grant write access to anyone. Grant EXECUTE only to those who intend to use Oracle Direct Connector for HDFS.

[Example 1–2](#) shows the SQL commands granting these privileges to `HDFSUSER`.

Example 1–2 Granting Users Access to Oracle Direct Connector for HDFS

```
CONNECT / AS sysdba;
CREATE USER hdfsuser IDENTIFIED BY password;
GRANT CREATE SESSION TO hdfsuser;
GRANT EXECUTE ON SYS.UTL_FILE TO hdfsuser;
GRANT READ, EXECUTE ON DIRECTORY hdfs_bin_path TO hdfsuser;
```

Oracle Loader for Hadoop Setup

Before installing Oracle Loader for Hadoop, verify that you have the required software.

Required Software

Oracle Loader for Hadoop requires the following software:

- A target database system running one of the following:
 - Oracle Database 10g release 2 (10.2.0.5) with required patch
 - Oracle Database 11g release 2 (11.2.0.2) with required patch
 - Oracle Database 11g release 2 (11.2.0.3)

Note: To use Oracle Loader for Hadoop with Oracle Database 10g release 2 (10.2.0.5) or Oracle Database 11g release 2 (11.2.0.2), you must first apply a one-off patch that addresses bug number 11897896. To access this patch, go to <http://support.oracle.com> and search for the bug number.

- Cludera's Distribution including Apache Hadoop version 3 (CDH3) or Apache Hadoop 0.20.2
- Hive 0.7.0 or 0.7.1, if using the `HiveToAvroInputFormat` class

Installing Oracle Loader for Hadoop

Oracle Loader for Hadoop is packaged with the Oracle Database 11g release 2 client libraries and Oracle Instant Client libraries for connecting to Oracle Database 10.2.0.5, 11.2.0.2, or 11.2.0.3.

To install Oracle Loader for Hadoop:

1. Unpack the content of the `oraloader-version.zip` archive into a directory on your Hadoop cluster.

A directory named `oraloader-version` is created with the following subdirectories:

- `doc`
 - `jlib`
 - `lib`
 - `examples`
2. Create a variable named `$OLH_HOME` and set it to the installation directory.
 3. Add `$OLH_HOME/jlib/*` to the `HADOOP_CLASSPATH` variable.

Oracle Data Integrator Application Adapter for Hadoop Setup

Installation requirements for Oracle Data Integrator (ODI) Application Adapter for Hadoop are provided in these topics:

- [System Requirements and Certifications](#)
- [Technology-Specific Requirements](#)

- [Location of ODI Application Adapter for Hadoop](#)
- [Setting Up the Topology](#)

System Requirements and Certifications

To use ODI Application Adapter for Hadoop, you must first have Oracle Data Integrator, which is licensed separately from Oracle Big Data Connectors. You can download ODI from the Oracle website at

<http://www.oracle.com/technetwork/middleware/data-integrator/downloads/index.html>

ODI Application Adapter for Hadoop requires a minimum version of Oracle Data Integrator 11.1.1.6.0.

Before performing any installation, read the system requirements and certification documentation to ensure that your environment meets the minimum installation requirements for the products that you are installing.

The list of supported platforms and versions is available on Oracle Technology Network:

<http://www.oracle.com/technetwork/middleware/data-integrator/overview/index.html>

Technology-Specific Requirements

The list of supported technologies and versions is available on Oracle Technical Network:

<http://www.oracle.com/technetwork/middleware/data-integrator/overview/index.html>

Location of ODI Application Adapter for Hadoop

ODI Application Adapter for Hadoop is available in the xml-reference directory of the Oracle Data Integrator Companion CD.

Setting Up the Topology

To set up the topology, see [Chapter 4, "Oracle Data Integrator Application Adapter for Hadoop."](#)

Oracle R Connector for Hadoop Setup

Oracle R Connector for Hadoop requires the installation of a software environment on the Hadoop side and on a client Linux system.

Installing the Software on Hadoop

Oracle Big Data Appliance supports Oracle R Connector for Hadoop without any additional software installation or configuration.

However, to use Oracle R Connector for Hadoop on any other Hadoop cluster, you must create the necessary environment.

Software Requirements for a Third-Party Hadoop Cluster

You must install several software components on a third-party Hadoop cluster to support Oracle R Connector for Hadoop.

Install these components on third-party servers:

- Cloudera's Distribution including Apache Hadoop version 3 (CDH3) or Apache Hadoop 0.20.2.
Complete the instructions provided by the distributor.
- Sqoop for the execution of functions that connect to Oracle Database. Oracle R Connector for Hadoop does not require Sqoop to install or load.
See "[Installing Sqoop on a Hadoop Cluster](#)" on page 1-8.
- Java Virtual Machine (JVM), preferably Java HotSpot Virtual Machine 6.
Complete the instructions provided at the download site at <http://www.oracle.com/technetwork/java/javase/downloads/index.html>
- R distribution 2.13.2 with all base libraries on all nodes in the Hadoop cluster.
Go to <http://www.r-project.org/> and follow the download and installation instructions.
- The ORCH package on each R engine, which must exist on every node of the Hadoop cluster.
See "[Installing the ORCH Package on a Hadoop Cluster](#)" on page 1-9.
- Open-source R `bitops` package installed on every node of the cluster, to support `orch.pack` and `orch.unpack`.

Installing Sqoop on a Hadoop Cluster

Sqoop provides a SQL-like interface to Hadoop, which is a Java-based environment. Oracle R Connector for Hadoop uses Sqoop for access to Oracle Database.

To install and configure Sqoop for use with Oracle Database:

1. Install Sqoop if it is not already installed on the server.
For Cloudera's Distribution including Apache Hadoop, see the Sqoop installation instructions in the *CDH Installation Guide* at <http://oracle.cloudera.com/>
2. Download the appropriate Java Database Connectivity (JDBC) driver for Oracle Database from Oracle Technology Network at <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>
3. Copy the driver JAR file to `$$SQOOP_HOME/lib`, which is a directory such as `/usr/lib/sqoop/lib`.
4. Provide Sqoop with the connection string to Oracle Database.

```
$ sqoop import --connect jdbc_connection_string
```


For example, `sqoop import --connect jdbc:oracle:thin@myhost:1521/orcl`.

Installing R on a Hadoop Cluster

You can download Oracle R Distribution 2.13.2 and get the installation instructions from the website at

<http://www.oracle.com/technetwork/indexes/downloads/r-distribution-1532464.html>

Installing the ORCH Package on a Hadoop Cluster

ORCH is the name of the Oracle R Connector for Hadoop package.

To install the ORCH package:

1. Set the environment variables for the supporting software:

```
$ setenv HADOOP_HOME /usr/lib/hadoop-0.2.0
$ setenv JAVA_HOME /usr/lib/jdk6
$ setenv R_HOME /usr/lib64/R
$ setenv SQOOP_HOME /usr/lib/sqoop
```

2. Unzip the downloaded file:

```
$ unzip orch-version.tgz.zip
Archive:  orch-0.1.8.tgz.zip
```

3. Open R and install the package:

```
> install.packages("/home/tmp/orch-0.1.8.tgz", repos=NULL)
Installing package(s) into ...
.
.
.
Hadoop 0.20.2-cdh3u3 is up
Sqoop 1.3.0-cdh3u3 is up
```

4. Alternatively, you can install the package from the Linux command line:

```
$ R CMD INSTALL orch-0.1.8.tgz
* installing *source* package 'ORCH' ...
** R
.
.
.
Hadoop 0.20.2-cdh3u3 is up
Sqoop 1.3.0-cdh3u3 is up

* DONE (ORCH)
```

Providing Remote Client Access to R Users

Whereas R users will run their programs as MapReduce jobs on the Hadoop cluster, they do not typically have individual accounts on that platform. Instead, an external Linux server provides remote access.

Software Requirements for Remote Client Access

To provide access to a Hadoop cluster to R users, install these components on a Linux server:

- The same version of Hadoop as your Hadoop cluster
- The same version of Sqoop as your Hadoop cluster

- The same version of the Java Development Kit (JDK) as your Hadoop cluster
- R distribution 2.13.2 with all base libraries
- ORCH R package

To provide access to database objects, you must have the Oracle Advanced Analytics option to Oracle Database. Then you can install this additional component on the Hadoop client:

- Oracle R Enterprise Client Packages

Configuring the Server as a Hadoop Client

You must install Hadoop on the client and minimally configure it for HDFS client use.

To install and configure Hadoop on the client system:

1. Install and configure CDH3 or Apache Hadoop 0.20.2 on the client system. This system can be the host for Oracle Database. If you are using Oracle Big Data Appliance, then complete the procedures for providing remote client access in the *Oracle Big Data Appliance Software User's Guide*. Otherwise, follow the installation instructions provided by the distributor (Cloudera or Apache).
2. Log in to the client system as an R user.
3. Open a Bash shell and enter this Hadoop file system command:

```
$HADOOP_HOME/bin/hadoop fs -ls /user
```
4. If you see a list of files, then you are done. If not, then ensure that the Hadoop cluster is up and running. If that does not fix the problem, then you must debug your client Hadoop installation.

Installing Sqoop on a Hadoop Client

Complete the same procedures on the client system for installing and configuring Sqoop as those provided in "[Installing Sqoop on a Hadoop Cluster](#)" on page 1-8.

Installing R on a Hadoop Client

You can download R 2.13.2 and get the installation instructions from the Oracle R Distribution website at

<http://oss.oracle.com/ORD/>

When you are done, ensure that users have the necessary permissions to connect to the Linux server and run R.

You may also want to install RStudio Server to facilitate access by R users. See the RStudio website at

<http://rstudio.org/>

Installing the ORCH Package on a Hadoop Client

Complete the procedures on the client system for installing ORCH as described in "[Installing the Software on Hadoop](#)" on page 1-7.

Installing the Oracle R Enterprise Client Packages (Optional)

To support access to Oracle Database, install the Oracle R Enterprise client packages. Without them, Oracle R Connector for Hadoop operates only with in-memory R

objects and local data files, and does not have access to Oracle Database or to the advanced statistical algorithms provided by Oracle R Enterprise.

See Also: *Oracle R Enterprise User's Guide* for information about installing R and Oracle R Enterprise

Oracle Direct Connector for Hadoop Distributed File System

This chapter describes how use Oracle Direct Connector for Hadoop Distributed File System (Oracle Direct Connector for HDFS) to facilitate data access between Hadoop Distributed File System (HDFS) and Oracle Database.

This chapter contains the following sections:

- [About Oracle Direct Connector for HDFS](#)
- [Creating an External Table for HDFS](#)
- [Publishing the HDFS Data Paths](#)
- [Querying Data in HDFS](#)

About Oracle Direct Connector for HDFS

Oracle Direct Connector for HDFS runs on the system where Oracle Database runs. It provides read access to HDFS from Oracle Database by using external tables.

An **external table** is an Oracle Database object that identifies the location of data outside of a database. Oracle Database accesses the data by using the metadata provided when the external table was created. By querying the external tables, users can access data stored in HDFS as if that data were stored in tables in a database. External tables are often used to stage data to be transformed during a database load.

These are a few ways that you can use Oracle Direct Connector for HDFS:

- To access any data stored in HDFS files
- To access comma-separated values (CSV) files and Oracle Data Pump files generated by Oracle Loader for Hadoop

CSV files contain plain text in delimited record format. The delimiter is typically a comma. Data Pump files are written in a proprietary, binary format.

- To load data extracted and transformed by Oracle Data Integrator

Oracle Direct Connector for HDFS uses the `ORACLE_LOADER` access driver.

Note: Oracle Direct Connector for HDFS requires a patch to Oracle Database before the connector can access Data Pump files produced by Oracle Loader for Hadoop. To download this patch, go to <http://support.oracle.com> and search for bug 14557588.

See Also:

- *Oracle Database Administrator's Guide* for information about external tables
- *Oracle Database Utilities* for more information about external tables, performance hints, and restrictions when you are using the ORACLE_LOADER access driver

Creating an External Table for HDFS

You create an external table for HDFS the same way as any other external table, except that you must specify this `PREPROCESSOR` clause in the SQL `CREATE TABLE` command:

```
PREPROCESSOR HDFS_BIN_PATH:hdfs_stream
```

`HDFS_BIN_PATH` is the name of the database directory that points to the `bin` subdirectory where Oracle Direct Connector for HDFS is installed. See ["Installing Oracle Direct Connector for HDFS"](#) on page 1-4.

To access Data Pump files, you must also specify this access parameter:

```
EXTERNAL VARIABLE DATA
```

Basic SQL Syntax for the External Table

Following is the basic SQL syntax for creating an external table for HDFS:

```
CREATE TABLE [schema.]table
  ( column datatype, ...
  )
  ORGANIZATION EXTERNAL
  (
    TYPE ORACLE_LOADER
    DEFAULT DIRECTORY directory
    ACCESS PARAMETERS
    ( PREPROCESSOR HDFS_BIN_PATH:hdfs_stream
      access_parameters...
    )
    LOCATION (file1,file2...)
  );
```

schema.table

Name of the external table to be created.

column

Name of the columns in the table.

datatype

Data type of the column, which is limited to the data types supported by ORACLE_LOADER. Oracle Database performs some data type conversions automatically.

directory

Name of the default directory object used by the external table for all input and output files, such as the location files, log files, and bad record files. Do not use the directory where the data is stored for these files.

access_parameters

Any additional subclauses in the ORACLE_LOADER `access_parameters` clause, such as record and field formatting.

file1, file2...

The names of the location files that identify the paths to the data in HDFS. For CSV content in HDFS, specify two or more location file names, because the degree of parallelism is limited by the number of location files. For Data Pump content in HDFS, specify one location file for each Data Pump file.

Oracle Direct Connector for HDFS creates these files in the default directory. If the files already exist, they are overwritten.

For other types of external tables, the location files contain the data, but Oracle Direct Connector for HDFS retains the data in HDFS.

Testing the External Table

After creating the external table, query it to verify that the preprocessor script is configured correctly:

```
SELECT count(*) FROM external_table;
```

If the query returns no rows and no errors, then you can continue.

External Table Example

[Example 2–1](#) creates an external table named SALES_HDFS_EXT_TAB in the SCOTT schema. The SALES_HDFS_EXT_TAB external table is created in a database directory named SALES_EXT_DIR. SCOTT must have read and write privileges on this directory.

To create the SALES_EXT_DIR database directory:

1. Create the file system directory:

```
$ mkdir /scratch/sales_ext_dir
$ chmod 664 /scratch/sales_ext_dir
```

2. Open a SQL command interface:

```
$ sqlplus / as sysdba
```

3. Create the database directory:

```
SQL> CREATE OR REPLACE DIRECTORY sales_ext_dir AS '/scratch/sales_ext_dir'
```

4. Grant read and write access to SCOTT:

```
SQL> GRANT READ, WRITE ON DIRECTORY sales_ext_dir TO scott;
```

Example 2–1 Defining an External Table for HDFS

```
CREATE TABLE "SCOTT"."SALES_HDFS_EXT_TAB"
( "PROD_ID"          NUMBER(6),
  "CUST_ID"          NUMBER,
  "TIME_ID"          DATE,
  "CHANNEL_ID"       CHAR(1),
  "PROMO_ID"         NUMBER(6),
  "QUANTITY_SOLD"    NUMBER(3),
  "AMOUNT_SOLD"      NUMBER(10,2)
)
ORGANIZATION EXTERNAL
( TYPE ORACLE_LOADER
  DEFAULT DIRECTORY "SALES_EXT_DIR"
  ACCESS PARAMETERS
  ( RECORDS DELIMITED BY NEWLINE
```

```

PREPROCESSOR HDFS_BIN_PATH:hdfs_stream
FIELDS TERMINATED BY ','
(
  "PROD_ID" DECIMAL EXTERNAL,
  .
  .
  "TIME_ID" CHAR DATE_FORMAT TIMESTAMP MASK "...",
  .
  .
)
)
LOCATION ( 'sales1','sale2','sales3')
);

```

Publishing the HDFS Data Paths

The previous procedure for creating an external table only created the metadata in Oracle Database. As mentioned earlier, the location files typically store the data values. In this case, however, the location files are empty. By executing the Oracle Direct Connector for HDFS `ExternalTable` command-line tool, you populate the location files with the Universal Resource Identifiers (URIs) of the data files in HDFS.

When users query the external table, the Oracle Direct Connector for HDFS preprocessor uses this information to locate the data in HDFS and stream it to the database.

ExternalTable Command

The `ExternalTable` command uses the values of several properties to populate the location files. You can specify these property values in an XML document or individually on the command line.

Altering HADOOP_CLASSPATH

Before issuing the `ExternalTable` command, alter the `HADOOP_CLASSPATH` environment variable to include these JAR files:

- `$ORACLE_HOME/jdbc/lib/ojdbc6.jar`: Required.
- `$ORACLE_HOME/jlib/oraclepki.jar`: Required only if you are using an Oracle wallet as an external password store.

See "[ExternalTable Command Example](#)" on page 2-5.

ExternalTable Command Syntax

This is the syntax of the `ExternalTable` command:

```

$HADOOP_HOME/bin/hadoop jar $DIRECTHDFS_HOME/jlib/orahdfs.jar
oracle.hadoop.hdfs.exttab.ExternalTable [-conf config_file | -D property=value]
-publish [--noexecute]

```

-conf *config_file*

Identifies the name of an XML configuration file containing the properties needed to populate the location files. See "[Creating a Configuration File](#)" on page 2-5.

-D *property=value*

Assigns a value to a specific property.

--noexecute

Generates an execution plan, but does not execute any of the actions.

ExternalTable Command Example

[Example 2-2](#) sets the HADOOP_CLASSPATH variable and publishes HDFS data paths to the external table created in [Example 2-1](#).

Example 2-2 Publishing HDFS Data Paths to an External Table

This example uses the Bash shell.

```
$ export HADOOP_CLASSPATH="$ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_
HOME/jlib/oraclepki.jar"

$ $HADOOP_HOME/bin/hadoop jar \
  $DIRECTHDFS_HOME/jlib/orahdfs.jar oracle.hadoop.hdfs.exttab.ExternalTable \
  -D oracle.hadoop.hdfs.exttab.tableName=SALES_HDFS_EXT_TAB \
  -D oracle.hadoop.hdfs.exttab.datasetPaths=hdfs:/user/scott/data/ \
  -D oracle.hadoop.hdfs.exttab.connection.url=jdbc:oracle:thin@myhost:1521/orcl \
  -D oracle.hadoop.hdfs.exttab.connection.user=scott -publish
```

In this example:

- \$HADOOP_HOME is an environment variable pointing to the Hadoop home directory.
- \$DIRECTHDFS_HOME is an environment variable pointing to the Oracle Direct Connector for HDFS installation directory.
- SALES_HDFS_EXT_TAB is the external table created in [Example 2-1](#).
- hdfs:/user/scott/data/ is the location of the HDFS data.
- @myhost:1521/orcl is the database connection string.

How to Publish Data Pump Files

When you are publishing Data Pump files:

- Set this property:


```
oracle.hadoop.hdfs.exttab.datapumpMode=true
```
- Ensure that the number of location files of the external table equals the number of files in the data set.

If the publish operation fails with an error message that indicates that the number of files in the data set doesn't match the number of location files, then alter the location clause of the external table to match the number of data set files listed in the error message and retry the publish operation.

Creating a Configuration File

A configuration file is an XML document with a very simple structure as follows:

```
<?xml version="1.0"?>
<configuration>
  <property>
    <name>property</name>
    <value>value</value>
  </property>
  .
  .
```

```
</configuration>
```

See "Configuration Properties" on page 2-6 for descriptions of these properties.

[Example 2-3](#) shows a configuration file.

Example 2-3 Configuration File for Oracle Direct Connector for HDFS

```
<?xml version="1.0"?>
<configuration>
  <property>
    <name>oracle.hadoop.hdfs.exctab.tableName</name>
    <value>SH.SALES_EXT_DIR</value>
  </property>
  <property>
    <name>oracle.hadoop.hdfs.exctab.datasetPaths</name>
    <value>/data/s1/*.csv,/data/s2/*.csv</value>
  </property>
  <property>
    <name>oracle.hadoop.hdfs.exctab.datasetCompressionCodec</name>
    <value>org.apache.hadoop.io.compress.DefaultCodec</value>
  </property>
  <property>
    <name>oracle.hadoop.hdfs.exctab.connection.url</name>
    <value>
      jdbc:oracle:thin:@example.com:1521/example.example.com
    </value>
  </property>
  <property>
    <name>oracle.hadoop.hdfs.exctab.connection.user</name>
    <value>SH</value>
  </property>
</configuration>
```

Configuration Properties

Following are the configuration properties used by the ExternalTable command to create links to the data files in HDFS. These properties are required:

- [oracle.hadoop.hdfs.exctab.connection.url](#)
- [oracle.hadoop.hdfs.exctab.datasetPaths](#)
- [oracle.hadoop.hdfs.exctab.tableName](#)
- [oracle.hadoop.hdfs.exctab.datapumpMode](#) for Data Pump files only

Property Descriptions

oracle.hadoop.hdfs.exctab.connection.tns_admin

File path to a directory containing Oracle Net configuration files, such as `sqlnet.ora` and `tnsnames.ora`. The value of the `TNS_ADMIN` environment variable is used for this property by default.

Define this property to use TNS entry names in database connect strings.

oracle.hadoop.hdfs.exctab.connection.tnsEntryName

A TNS entry name defined in the `tnsnames.ora` file. This property is used with [oracle.hadoop.hdfs.exctab.connection.tns_admin](#).

oracle.hadoop.hdfs.exttab.connection.url

The URL of the database connection string. This property overrides all other connection properties. The connecting database user must have the privileges described in "[Granting User Access to Oracle Direct Connector for HDFS](#)" on page 1-5. Required.

Using a Wallet

If you are using an Oracle wallet as an external password store, then the property value must have this form:

```
jdbc:oracle:thin:@db_connect_string
```

The `db_connect_string` must exactly match the credential in the wallet.

This example uses Oracle Net Services syntax:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=myhost)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=my_db_service_name)))
```

The next example uses a TNSNAMES entry:

```
jdbc:oracle:thin:@my_tns_entry
```

See [oracle.hadoop.hdfs.exttab.connection.wallet_location](#). For information about wallets, see the *Oracle Database Advanced Security Administrator's Guide*.

Not Using a Wallet

If you are not using an Oracle wallet, then use one of the following URL connection styles:

- Thin Connection:

```
jdbc:oracle:thin:@//myhost:1521/my_db_service_name
```

- Oracle Net Services:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=myhost)(PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=my_db_service_name)))
```

- TNS Entry Name:

```
jdbc:oracle:thin:@myTNSEntryName
```

This parameter is required when you are not using an Oracle wallet for connections:

- [oracle.hadoop.hdfs.exttab.connection.user](#)

oracle.hadoop.hdfs.exttab.connection.user

An Oracle Database user name.

oracle.hadoop.hdfs.exttab.connection.wallet_location

File path to an Oracle wallet where the connection information is stored. When you are using a wallet as an external password store, set the following properties:

For a URL connection:

- `oracle.hadoop.hdfs.exttab.connection.wallet_location`
- [oracle.hadoop.hdfs.exttab.connection.url](#)

For TNS names:

- `oracle.hadoop.hdfs.exctab.connection.wallet_location`
- [oracle.hadoop.hdfs.exctab.connection.tns_admin](#)
- [oracle.hadoop.hdfs.exctab.connection.tnsEntryName](#)

oracle.hadoop.hdfs.exctab.datapumpMode

Indicates that Data Pump files are being published. Set to `true` when you are publishing Data Pump files.

oracle.hadoop.hdfs.exctab.datasetCompressionCodec

The class name of the compression codec that implements the `org.apache.hadoop.io.compress.CompressionCodec` interface. The Decompressor class from the codec is used by the preprocessor script to decompress data for the external table. This codec applies to the entire data set.

Specify this property if the data set contains compressed files.

oracle.hadoop.hdfs.exctab.datasetPathFilter

The class name of a path filter that implements the `org.apache.hadoop.fs.PathFilter` interface. The paths in the data set are selected only if this filter class accepts them.

oracle.hadoop.hdfs.exctab.datasetPaths

A comma-separated list of fully qualified HDFS paths. This parameter enables you to restrict the input by using special pattern-matching characters in the path specification. See [Table 2–1](#). Required.

For example, to select all files in `/data/s2/`, and only the CSV files in `/data/s7/`, `/data/s8/`, and `/data/s9/`, enter this expression:

```
/data/s2/,/data/s[7-9]/*.csv
```

The external table accesses the data contained in all listed files and all files in listed directories. These files compose a single data set.

The data set can contain compressed files or uncompressed files, but not both.

Table 2–1 *Pattern-Matching Characters*

Character	Description
<code>?</code>	Matches any single character
<code>*</code>	Matches zero or more characters
<code>[abc]</code>	Matches a single character from the character set $\{a, b, c\}$
<code>[a-b]</code>	Matches a single character from the character range $\{a\dots b\}$. The character a must be less than or equal to b .
<code>[^a]</code>	Matches a single character that is not from character set or range $\{a\}$. The carat (<code>^</code>) must immediately follow the left bracket.
<code>\c</code>	Removes any special meaning of character c . The backslash is the escape character.
<code>{ab\,cd}</code>	Matches a string from the string set $\{ab, cd\}$. Precede the comma with an escape character (<code>\</code>) to remove the meaning of the comma as a path separator.
<code>{ab\,c{de\,fh}}</code>	Matches a string from the string set $\{ab, cde, cfh\}$. Precede the comma with an escape character (<code>\</code>) to remove the meaning of the comma as a path separator.

oracle.hadoop.hdfs.exttab.tableName

Schema-qualified name of the external table in the format

schemaName.tableName

See ["Creating an External Table for HDFS"](#) on page 2-2 for information about creating an external table for Oracle Direct Connector for HDFS. Required.

Querying Data in HDFS

Parallel processing is extremely important when you are working with large volumes of data. When you use external tables, always enable parallel query with this SQL command:

```
ALTER SESSION ENABLE PARALLEL QUERY;
```

Before loading the data into an Oracle database from the external files created by Oracle Direct Connector for HDFS, enable parallel DDL:

```
ALTER SESSION ENABLE PARALLEL DDL;
```

Before inserting data into an existing database table, enable parallel DML with this SQL command:

```
ALTER SESSION ENABLE PARALLEL DML;
```

Hints such as APPEND and PQ_DISTRIBUTE also improve performance when you are inserting data.

Oracle Loader for Hadoop

This chapter explains how to use Oracle Loader for Hadoop to copy data from Hadoop files into tables in an Oracle database. It contains the following sections:

- [What Is Oracle Loader for Hadoop?](#)
- [Using Oracle Loader for Hadoop](#)
- [Output Modes During OraLoader Invocation](#)
- [Balancing Loads When Loading Data into Partitioned Tables](#)
- [Primary Configuration Properties for the Load Balancing Feature](#)
- [OraLoader Configuration Properties](#)
- [Example of Using Oracle Loader for Hadoop](#)
- [Target Table Characteristics](#)
- [Loader Map XML Schema Definition](#)
- [OraLoader for Hadoop Configuration Properties](#)
- [Third-Party Licenses for Bundled Software](#)

What Is Oracle Loader for Hadoop?

Oracle Loader for Hadoop is an efficient and high-performance loader for fast movement of data from a Hadoop cluster into a table in an Oracle database. Oracle Loader for Hadoop prepartitions the data if necessary and transforms it into a database-ready format. It optionally sorts records by primary key before loading the data or creating output files. Oracle Loader for Hadoop is a MapReduce application that is invoked as a command-line utility.

Note: Partitioning is a database feature for managing and efficiently querying very large tables. It provides a way to decompose a large table into smaller and more manageable pieces called partitions, in a manner entirely transparent to applications. For more information about partitioning, see the *Oracle Database VLDB and Partitioning Guide*.

After the prepartitioning and transforming steps, there are two modes for loading the data into an Oracle database from a Hadoop cluster:

- Online database mode: The data is loaded into the database using either a JDBC output format or an OCI Direct Path output format. The OCI Direct Path output

format performs a high-performance direct path load of the target table. The JDBC output format performs a conventional path load. For more information about these online load methods, including restrictions for the OCI Direct Path output format, see ["Output Modes During OraLoader Invocation"](#) on page 3-7.

- **Offline database mode:** The reducer nodes create binary or text format output files. The Oracle Data Pump output format creates binary format files that are ready to be loaded into an Oracle database using an external table and the `ORACLE_DATAPUMP` access driver. The Delimited Text output format creates text files in delimited record format. (This is usually called comma-separated values (CSV) format when the delimiter is a comma.) These text files are ready to be loaded into an Oracle database using an external table and the `ORACLE_LOADER` access driver. The files can also be loaded using the `SQL*Loader` utility.

Using Oracle Loader for Hadoop

This section describes the following steps for using Oracle Loader for Hadoop:

1. [Implementing InputFormat](#)
2. [Creating the loaderMap Document](#)
3. [Accessing Table Metadata](#)
4. [Invoking OraLoader](#)
5. [Loading Files Into an Oracle Database \(Offline Loads Only\)](#)

See ["Oracle Loader for Hadoop Setup"](#) on page 1-6 for installation instructions.

Implementing InputFormat

Oracle Loader for Hadoop is a MapReduce application that gets its input from an `org.apache.hadoop.mapreduce.RecordReader` implementation as provided by the `org.apache.hadoop.mapreduce.InputFormat` class that is specified in the `mapreduce.inputformat.class` configuration property. Oracle Loader for Hadoop requires that the `RecordReader` return an Avro `IndexedRecord` input object from the `getCurrentKey()` method. The method signature should be:

```
public org.apache.avro.generic.IndexedRecord getCurrentKey()
throws IOException, InterruptedException;
```

Oracle Loader for Hadoop uses the schema of the `IndexedRecord` input object to discover the names of the input fields and map them to the columns of the table to load. This mapping is discussed in more detail in the following sections.

Oracle Loader for Hadoop comes with two built-in input formats; it also provides the source code for two `InputFormat` examples. The example source code is located in the `jsrc/` directory. [Table 3-1](#) lists the class names for all these input formats, along with the types of input they handle and how they generate the Avro schema field names. (Understanding how an `InputFormat` generates field names is critical to getting the data loaded into the target table.)

The built-in input format classes are described in the next subsections. For the `InputFormat` examples, consult the source code and Javadoc for these classes for more information.

Table 3–1 InputFormat Classes, Types, and Field Names

Class	Input Type	Avro Schema Field Names
oracle.hadoop.loader.lib.input.HiveToAvroInputFormat	Hive table sources	Hive table's column names (upper cased)
oracle.hadoop.loader.lib.input.DelimitedTextInputFormat	Delimited text files	Comma-separated list from the property <code>oracle.hadoop.loader.input.fieldNames</code> (or F0, F1, ... if property is not defined)
oracle.hadoop.loader.examples.CSVInputFormat	Simple, delimited text files	F0, F1,...
oracle.hadoop.loader.examples.AvroInputFormat	Binary format Avro record files	Field names from input files' Avro schemas

HiveToAvroInputFormat

This class presents an input format that reads data from a Hive table. It requires that the Hive database and table names be specified using the following configuration properties:

- `oracle.hadoop.loader.input.hive.tableName`
- `oracle.hadoop.loader.input.hive.databaseName`

`HiveToAvroInputFormat` contacts the `HiveMetaStoreClient` to retrieve information about the table's columns, location, `InputFormat`, `SerDe`, and so on. Depending on the way in which Hive was configured, additional Hive-specific properties must be set (such as `hive.metastore.uris` and `hive.metastore.local`).

Note that Oracle Loader for Hadoop does not currently support Hive partitioned tables.

`HiveToAvroInputFormat` imports the entire table (all the files in the Hive table's directory). All other (file-based) input formats discussed in this document allow "globbing" (that is, appending wildcard patterns to the input directories to restrict the input).

Rows in the Hive table are transformed into Avro records whose field names are the Hive table's column names in uppercase. This makes it more likely that the fields will coincide with the database column names. See "[Creating the loaderMap Document](#)" on page 3-4.

DelimitedTextInputFormat

This is an `InputFormat` class for delimited text files, such as comma-separated values files or tab-separated values files. `DelimitedTextInputFormat` requires that records be separated by newline characters and that fields be delimited using single-character markers.

The `DelimitedTextInputFormat` class is meant to emulate the "terminated by **t** [optionally enclosed by **ie** [and **te**]]" behavior of `SQL*Loader`. The **t** is the field terminator, **ie** is the initial field encloser, and **te** is the trailing field encloser.

`DelimitedTextInputFormat` uses a parser based on the following grammar:

- `Line = Token t Line | Token\n`
- `Token = EnclosedToken | UnenclosedToken`

- EnclosedToken = (white-space)* ie [(non-te)* te te]* (non-te)* te (white-space)*
- UnenclosedToken = (white-space)* (non-t)*
- white-space = {c | Character.isWhitespace(c) and c!=t}

Any trailing field enclosure character contained inside an enclosed token must be encoded by "doubling it up" (that is, printing it twice).

White space around enclosed tokens is discarded. For unenclosed tokens the leading white space is discarded, but not the trailing white space (if any).

Any empty-string token (either enclosed or unenclosed) is replaced with a null.

This implementation allows custom enclosers and terminator characters (see [Table 3-2](#)), but hard codes the record terminator (to newline) and white space (to Java's `Character.isWhitespace()`). The enclosers must be different from the terminator character and white spaces (but can be equal to each other). The terminator can be a white space (but that value is removed from the class of white space characters).

[Table 3-2](#) describes the delimiters available for `DelimitedTextInputFormat`. In the table, HHHH is a big-endian hexadecimal representation of the character in UTF-16.

Table 3-2 Delimiters for DelimitedTextInputFormat

Delimiter Type	Property	Possible Values	Default
Field terminator	<code>oracle.hadoop.loader.input.fieldTerminator</code>	<ul style="list-style-type: none"> ■ one character ■ <code>\uHHHH</code> 	, (comma)
Initial field enclosure	<code>oracle.hadoop.loader.input.initialFieldEncloser</code>	<ul style="list-style-type: none"> ■ one character ■ <code>\uHHHH</code> ■ nothing 	No default
Trailing field enclosure	<code>oracle.hadoop.loader.input.trailingFieldEncloser</code>	<ul style="list-style-type: none"> ■ one character ■ <code>\uHHHH</code> ■ nothing 	No default

The field enclosers must be either both set or both not set. If they are not set, then the `EnclosedToken` non-terminal element is essentially removed from the grammar of the parser that `DelimitedTextInputFormat` uses. If field enclosers are set, then the parser attempts to read each field as an `EnclosedToken` first before reading it as an `UnenclosedToken`. Note that if the initial field enclosure is set, then the trailing field enclosure must also be set, even when the two are the same.

`DelimitedTextInputFormat` reads the field names as a comma-separated list from the configuration property `oracle.hadoop.loader.input.fieldNames`. If parsing a line results in more tokens (fields) than field names, the extra tokens are discarded. If there are fewer tokens than field names, the missing trailing tokens are set to null.

If the `oracle.hadoop.loader.input.fieldNames` property is not set, then the `DelimitedTextInputFormat`'s `RecordReader` uses `F0, F1, ... Fn` as field names (where `n` is the largest number of tokens encountered by that `RecordReader` in any line so far).

Creating the loaderMap Document

Oracle Loader for Hadoop loads data into a single database table. This table is referred to as the target table. You can use the following ways to specify the target table, the columns to load, and how input fields are mapped to database columns:

- To indicate that all columns of the database table will be loaded and that the names of the input fields exactly match the database column names, use the configuration property `oracle.hadoop.loader.targetTable`. It allows you to define a schema-qualified name for the target load table. For each database column, the loader uses the column name to discover an input field with the same name. The value of the field is then loaded into the column.
- To load a subset of the target table columns or to create explicit mappings when the input field names are not exactly the same as the database column names, create a `loaderMap` document to specify the target table, columns, and how the input fields should be mapped to the database columns. The location of the `loaderMap` document is specified using the `oracle.hadoop.loader.loaderMapFile` configuration property.

See Also:

- ["Target Table Characteristics"](#) on page 3-21
- ["Loader Map XML Schema Definition"](#) on page 3-22 for information about the content of the XML schema definition (XSD) document

Example loaderMap Document

The following example `loaderMap` document specifies a list of columns in the `HR.EMPLOYEES` table that should be loaded. It includes a mapping of input data field names to table column names. It also specifies the format of input data that should be used for that column.

```
<?xml version="1.0" encoding="UTF-8"?>
<LOADER_MAP>
<SCHEMA>HR</SCHEMA>
<TABLE>EMPLOYEES</TABLE>
<COLUMN field="empId">EMPLOYEE_ID</COLUMN>
<COLUMN field="lastName">LAST_NAME</COLUMN>
<COLUMN field="email">EMAIL</COLUMN>
<COLUMN field="hireDate" format="MM-dd-yyyy">HIRE_DATE</COLUMN>
<COLUMN field="jobId">JOB_ID</COLUMN>
</LOADER_MAP>
```

Note: If all the columns in the target table are used for loading, and if the input data field names in the `IndexedRecord` input object match the column names exactly, then the `loaderMap` file is not needed unless a table column is a `DATE` data type. Input fields mapped to `DATE` columns are parsed using the default Java date format. If the input is in a different format, then you must create a `loaderMap` document and use the `format` attribute to specify the Java date format string to use when parsing input values.

Accessing Table Metadata

Oracle Loader for Hadoop uses table metadata from Oracle Database to control the execution of a loader job. The loader automatically fetches the metadata whenever a JDBC connection can be established. Sometimes it may be impossible for the loader job to access the database. For example, the Hadoop cluster may be on a different network than the database. In this case, the `OraLoaderMetadata` utility program is used to extract table metadata from the database into an XML document. The metadata

document is then transferred to the Hadoop cluster. The configuration property `oracle.hadoop.loader.tableMetadataFile` is used to specify the location of the metadata document. When the loader job runs, it accesses this document to discover all necessary metadata information about the target table.

Running the OraLoaderMetadata Utility

To run the `OraLoaderMetadata` Java utility, add the following JAR files to the `CLASSPATH` variable:

- `$OLH_HOME/jlib/oraloader.jar`
- `$OLH_HOME/jlib/ojdbc6.jar`
- `$OLH_HOME/jlib/oraclepki.jar`

Note: The `oraclepki.jar` library is required only if you are connecting to the database using credentials stored in an Oracle wallet.

Then run the following command:

```
java oracle.hadoop.loader.metadata.OraLoaderMetadata \  
-user <username> -connection_url <connection URL> [-schema <schemaName>] \  
-table <tableName> -output <output filename>
```

OraLoaderMetadata Parameters

The `OraLoaderMetadata` utility accepts the following parameters:

- `-user` is the Oracle database user name. The user is prompted for the password.
- `-connection_url` is the connection URL to connect to Oracle Database.
- `-schema` is the name of the schema containing the target table. If this parameter is not specified, then the target table is assumed to be in the user schema specified in the connect URL.
- `-table` is the name of the target table.
- `-output` is the output file name used to store the metadata document.

Invoking OraLoader

`OraLoader` is a Hadoop job that you execute using the standard Hadoop tools. `OraLoader` implements the `org.apache.hadoop.util.Tool` interface and follows the standard Hadoop methods for building MapReduce applications. `OraLoader` performs the following actions:

1. Reads and verifies input configuration parameters.
2. Retrieves and verifies table and column metadata information for the target table. Metadata is retrieved from the database whenever a JDBC connection can be made. Otherwise, the loader looks for metadata stored in the location specified by the `oracle.hadoop.loader.tableMetadataFile` property.
3. Prepares internal configuration information for the MapReduce tasks of `OraLoader` and stores table metadata information and dependent Java libraries in the distributed cache so that they are available to the map and reduce tasks throughout the cluster.
4. Submits the MapReduce job to Hadoop.

5. Consolidates reporting information from individual tasks to create a common log file for the job after the map and reduce tasks are complete. The log file is written to the job output directory and is named `oraloader-report.txt`.

OraLoader is invoked from the command line and accepts any of the generic command-line options. The following is an example invocation:

```
HADOOP_CLASSPATH="$HADOOP_CLASSPATH:$OLH_HOME/jlib/*"
```

```
bin/hadoop jar $OLH_HOME/jlib/oraloader.jar oracle.hadoop.loader.OraLoader \
-conf MyConf.xml
```

When using one of the example input formats shown in [Table 3-1](#), you must include the `-libjars` option in the OraLoader invocation:

```
bin/hadoop jar $OLH_HOME/jlib/oraloader.jar oracle.hadoop.loader.OraLoader \
-conf MyConf.xml -libjars $OLH_HOME/jlib/oraloader-examples.jar
```

See Also:

- The Apache Hadoop documentation for information about where to find the Hadoop executable and the setting for the `HADOOP_CLASSPATH` variable
- Javadoc for the generic options, which is located at the following Apache site:
<http://hadoop.apache.org/common/docs/r0.20.2/api/org/apache/hadoop/util/GenericOptionsParser.html>

Loading Files Into an Oracle Database (Offline Loads Only)

For offline loads, Oracle Loader for Hadoop produces files that must be copied to the database server and loaded into an Oracle database. The following section describes the available offline load method.

Loading From Delimited Text Files Into an Oracle Database

After you copy the delimited text files to the database system, use the generated control files to invoke `SQL*Loader` and load the data from the delimited text files into the database. Alternatively, you can use the generated SQL scripts to perform external table loads into the database. See "[Delimited Text Output](#)" on page 3-9.

Output Modes During OraLoader Invocation

This section describes the following output options:

- [JDBC Output](#)
- [Oracle OCI Direct Path Output](#)
- [Delimited Text Output](#)
- [Oracle Data Pump Output](#)

JDBC Output

JDBC is an output option in online database mode. The output records of the loader job are loaded directly into the target table by map or reduce tasks as part of the OraLoader process. There is no need to execute additional steps to load the data. A JDBC connection between the Hadoop system and Oracle Database is required for this output option.

The JDBC output option uses standard JDBC batching to increase performance and efficiency. If an error occurs during batch execution (for example a constraint is violated), the JDBC driver stops execution at the first error. Thus, if there are 100 rows in a batch and the tenth row causes an error then nine rows are inserted and 91 rows are not. Moreover, the JDBC driver does not provide information to identify which row caused the error. In this case, Oracle Loader for Hadoop does not know the insert status for any of the rows in the batch. It counts all rows in the batch as "in question" and continues loading the next batch. A load report is produced at the end of the job that details the number of batch errors incurred and the number of rows whose insert status is in question. One way to handle this problem is by using a unique key on the data. After the data is loaded, the key can be enabled and used to discover missing key values. The missing rows must be located in the input data and reloaded after it has been determined why they failed to load.

To select the JDBC output format, set the following Hadoop property:

```
<property>
  <name>mapreduce.outputformat.class</name>
  <value>oracle.hadoop.loader.lib.output.JDBCOutputFormat</value>
</property>
```

The relevant property for configuring JDBC output is `oracle.hadoop.loader.jdbc.defaultExecuteBatch`, which controls the size of the batch.

Oracle OCI Direct Path Output

The Oracle OCI Direct Path output format is available in online database mode. This output format uses the OCI Direct Path interface to load rows into the target table. Parallel direct path load is possible because each reducer loads into a distinct database partition.

To select the Oracle OCI Direct Path output format, set the following Hadoop property:

```
<property>
  <name>mapreduce.outputformat.class</name>
  <value>oracle.hadoop.loader.lib.output.OCIOutputFormat</value>
</property>
```

The size of the direct path stream buffer can be controlled using the following property:

```
<property>
  <name>oracle.hadoop.loader.output.dirpathBufsize</name>
  <value>131072</value>
  <description>
    This property is used to set the size, in bytes, of the direct path
    stream buffer for OCIOutputFormat. If needed, values are rounded
    up to the next nearest multiple of 8k.
  </description>
</property>
```

The Oracle OCI Direct Path output format has the following restrictions:

- It is only available on a Linux x86.64 platform.
- The load target table must be partitioned.
- The number of reducers must be greater than zero.

- OCI Direct Path output cannot load a composite interval partitioned table where the subpartition key contains a CHAR, VARCHAR2, NCHAR, or NVARCHAR2 column. The loader checks for this condition and stops with an error if the target load table meets this condition. Composite interval partitions where the subpartition key does not contain a character type column are supported.

The Oracle OCI Direct Path output format requires the following configuration steps. These steps enable the loader to locate the C shared libraries that implement the output format. These libraries are automatically distributed to compute nodes using the Hadoop Distributed Cache mechanism.

1. Create the environment variable `JAVA_LIBRARY_PATH` to point to the directory `$OLH_HOME/lib`. This environment variable is required only on the node where the job is submitted. The `$HADOOP_HOME/bin/hadoop` command in CDH automatically injects this variable value into the Java system property `java.library.path` when the job is created. For the Apache Hadoop distribution, you must edit the `$HADOOP_HOME/bin/hadoop` command so that it concatenates new values to an existing value. The Apache `hadoop` command begins with an empty `JAVA_LIBRARY_PATH` value and does not import a value from the environment.
2. Add `$OLH_HOME/lib` to the `LD_LIBRARY_PATH` variable on the client where the loader job is submitted.

Delimited Text Output

Delimited text is an output option in offline database mode. Comma-separated value (CSV) format files, or other delimited text files, are generated by map or reduce tasks. These files are then loaded into the target table by using either SQL*Loader or external tables.

To select the Delimited Text output format, set the following Hadoop property:

```
<property>
  <name>mapreduce.outputformat.class</name>
  <value>oracle.hadoop.loader.lib.output.DelimitedTextOutputFormat</value>
</property>
```

Each output task generates a delimited text format file and a SQL*Loader control file (see [Example 3-1](#)). If the table is not partitioned or if `oracle.hadoop.loader.loadByPartition=false`, then a single SQL*Loader control file is generated for the entire job. Additionally, a single SQL script is generated to load the delimited text file into the target table.

Delimited text files have the following template:

```
oraloader-${taskId}-csv-${partitionId}.dat
```

SQL*Loader control file names have the following template:

```
oraloader-${taskId}-csv-${partitionId}.ctl
```

Definitions of the template parameters are as follows:

`${taskId}`: mapper (reducer) Id

`${partitionId}`: Partition identifier

If the table is not partitioned or if `oracle.hadoop.loader.loadByPartition=false`, then a single SQL*Loader control file is generated: `oraloader-csv.ctl`.

The SQL script for loading into an external table is called `oraloader-csv.sql`.

The formatting of records and fields in the delimited text file is controlled by the following properties:

- `oracle.hadoop.loader.output.fieldTerminator`: Identifies a single character that delimits the fields.
- `oracle.hadoop.loader.output.initialFieldEncloser`: Identifies the beginning of a field. It indicates that the fields are always enclosed between this character and the `trailingFieldEncloser` character.
- `oracle.hadoop.loader.output.trailingFieldEncloser`: Identifies the end of a field that begins with the `initialFieldEncloser` character. If this property is not set, then the value of the `initialFieldEncloser` is used as the field terminator.
- `oracle.hadoop.loader.output.escapeEnclosers`: Identifies the character that is used to escape embedded trailing field encloser characters.

[Example 3-1](#) shows a sample SQL*Loader control file that might be generated by an output task.

Example 3-1 Sample SQL*Loader Control File

```
LOAD DATA CHARACTERSET AL32UTF8
INFILE 'oraloader-csv-1-0.dat'
BADFILE 'oraloader-csv-1-0.bad'
DISCARDFILE 'oraloader-csv-1-0.dsc'
INTO TABLE "SCOTT"."CSV_PART" PARTITION(10) APPEND
FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '''
(
  "ID"          DECIMAL EXTERNAL,
  "NAME"       CHAR,
  "DOB"        DATE 'SYYYY-MM-DD HH24:MI:SS'
)
```

Oracle Data Pump Output

The Oracle Data Pump output format is available in offline database mode. The loader produces binary format files that can be loaded into the target table using an external table and the `ORACLE_DATAPUMP` access driver. The output files must be copied from the HDFS file system to a local file system that is accessible to Oracle Database.

To select the Oracle Data Pump output format, set the following Hadoop property:

```
<property>
  <name>mapreduce.outputformat.class</name>
  <value>oracle.hadoop.loader.lib.output.DataPumpOutputFormat</value>
</property>
```

Oracle Data Pump output file names have the following template:

```
oraloader-${taskId}-dp-${partitionId}.dat
```

Oracle Loader for Hadoop also produces a SQL file that contains commands to perform the following tasks:

1. Create an external table definition using the `ORACLE_DATAPUMP` access driver. The binary format Oracle Data Pump output files are listed in the `LOCATION` clause of the external table.
2. Create a directory object that is used by the external table. This command must be uncommented before it can be used. To specify the directory name that is produced in the SQL file, set the following property:

```

<property>
  <name>oracle.hadoop.loader.extTabDirectoryName</name>
  <value>OLH_EXTTAB_DIR</value>
  <description>
    The name of the Oracle directory object for the external table's
    LOCATION data files. This property applies only to the CSV and
    DataPump output formats.
  </description>
</property>

```

3. Insert the rows from the external table into the target table. This command must be uncommented before it can be used.

See Also:

- *Oracle Database Administrator's Guide* for more information about creating and managing external tables
- *Oracle Database Utilities* for more information about the ORACLE_DATAPUMP access driver

Balancing Loads When Loading Data into Partitioned Tables

To balance loads across reducers when data is loaded into a partitioned database table, use the sampling feature of Oracle Loader for Hadoop.

The execution time of a reducer is usually proportional to the number of records that it processes—the more records, the longer the execution time. When the sampling feature is disabled, all records from a given database partition are sent to one reducer. This can result in unbalanced reducer loads because some database partitions may have more records than others. Because the execution time of a Hadoop job is the execution time of its slowest reducer, unbalanced reducer loads can slow down the entire job.

Although hashing records uniformly across reducers can generate balanced reducer loads, it does not necessarily group records by database partition before inserting them into the database.

The sampling feature of Oracle Loader for Hadoop generates an efficient MapReduce partitioning scheme that groups records by database partition while also balancing reducer load.

Using the Sampling Feature

To enable the sampling feature, set the configuration property `oracle.hadoop.loader.sampler.enableSampling` to `true`.

Even if the `enableSampling` property is set to `true`, the loader automatically disables the sampling feature if sampling is not necessary or if the loader determines that a good sample cannot be made. For example, sampling is automatically disabled if the table is not partitioned, the number of reducer tasks is less than two, or there is too little input data to compute a good load balance. In those cases, the loader returns an informational message.

Note: The sampler is multithreaded, and each sampler thread instantiates its own copy of the supplied `InputFormat` class. Any new `InputFormat` implementations provided to Oracle Loader for Hadoop should ensure that data structures that are static and mutable are synchronized for multiple thread access.

It is possible for the sampler to return an out-of-memory error on the client node where the loader job is submitted. This can occur when the input splits returned by the `InputFormat` do not fit in memory.

The following are possible solutions to this problem:

- Increase the heap size of the JVM where the job is submitted.
- Adjust the following properties:

```
oracle.hadoop.loader.sampler.hintMaxSplitSize
oracle.hadoop.loader.sampler.hintNumMapTasks
```

See "[OraLoader for Hadoop Configuration Properties](#)" on page 3-23 for descriptions of these properties.

Tuning Load Balancing and Sampling Behavior

Oracle Loader for Hadoop provides properties that you can use to tune load balancing and sampling behavior. These properties are summarized in [Table 3–3](#) on page 3-14.

Properties to Tune Load Balancing

The goal of load balancing is to generate a MapReduce partitioning scheme that assigns approximately the same amount of work to all reducers. This scheme is used in the partitioning step during Oracle Loader for Hadoop job execution.

Two properties control the quality of load balancing: `maxLoadFactor` and `loadCI`. The sampler uses the expected reducer load factor to evaluate the quality of its partitioning scheme. **Load factor** is a metric that indicates how much a reducer's load deviates from a perfectly balanced reducer load. A load factor of one indicates a perfectly balanced load (no overload).

Small load factors indicate better load balancing. The `maxLoadFactor` default of 0.05 means that no reducer is ever overloaded by more than 5%. The sampler guarantees this `maxLoadFactor` with a statistical confidence level determined by the value of `loadCI`. The default value of `loadCI` is 0.95, which means that any reducer's load factor exceeds `maxLoadFactor` in only 5% of the cases.

There is a trade-off between the execution time of the sampler and the quality of load balancing. Lower values of `maxLoadFactor` and higher values of `loadCI` result in more balanced reducer loads at the expense of longer sampling times. The default values of `maxLoadFactor=0.05` and `loadCI=0.95` provide a good trade-off between load balancing quality and execution time.

Properties to Tune Sampling Behavior

By default, the sampler runs until it collects just enough samples to generate a partitioning scheme that satisfies the `maxLoadFactor` and `loadCI` criteria.

However, you can limit the sampler's running time by using the `maxSamplesPct` property, which specifies the maximum number of records that the sampler should sample before stopping.

Does Oracle Loader for Hadoop Always Use the Sampler's Partitioning Scheme?

Oracle Loader for Hadoop uses the generated partitioning scheme only if sampling is successful. A sampling is successful if it generates a partitioning scheme with a maximum reducer load factor of $(1 + \text{maxLoadFactor})$ guaranteed at a statistical confidence level of `loadCI`. The default values of `maxLoadFactor`, `loadCI`, and `maxSamplesPct` allow the sampler to successfully generate high-quality partitioning schemes for a variety of different input data distributions. However, in some cases the sampler might be unsuccessful in generating a partitioning scheme that satisfies these constraints (for example, if the constraints are too rigid or if the number of samples it requires exceeds the user-specified maximum of `maxSamplesPct`). In such cases, Oracle Loader for Hadoop prints a log message saying that there were not enough samples. It then defaults to partitioning records by database partition and provides no load balancing guarantees (as described in ["Tuning Load Balancing and Sampling Behavior"](#) on page 3-12).

An alternative approach would be to reset the configuration properties to less rigid values. You can do this either by increasing `maxSamplesPct` or by decreasing `maxLoadFactor` or `loadCI`, or both.

What Happens When a Sampling Feature Property Has an Invalid Value?

If any configuration properties of the sampling feature are set to values outside the accepted range, an exception is not returned. Instead, the sampler prints a warning message, resets the property to its default value, and continues executing.

Primary Configuration Properties for the Load Balancing Feature

[Table 3–3](#) describes the primary properties available to tune sampling behavior. See ["OraLoader for Hadoop Configuration Properties"](#) on page 3-23 for a complete list of properties.

Table 3–3 Configuration Properties for the Oracle Loader for Hadoop Sampling Feature

Property Name	Attributes and Description
<code>oracle.hadoop.loader.sampler.maxSamplesPct</code>	<p>Type: Float</p> <p>Default: 0.01</p> <p>Accepted Range: [0, 1]</p> <p>A value of ≤ 0 disables this property.</p> <p>Description: The maximum sample size as a percentage of the number of records in the input data. A value of 0.05 indicates that the sampler never samples more than 5% of the total number of records. The sampler may collect fewer samples than this amount.</p>
<code>oracle.hadoop.loader.sampler.maxLoadFactor</code>	<p>Type: Float</p> <p>Default: 0.05</p> <p>Accepted Range: ≥ 0</p> <p>A value of ≤ 0 resets the property to the default.</p> <p>Description: Maximum acceptable load factor for reducer workload.</p>
<code>oracle.hadoop.loader.sampler.loadCI</code>	<p>Type: Float</p> <p>Default: 0.95</p> <p>Accepted Range: ≥ 0.5 and < 1</p> <p>Recommended values are ≥ 0.9.</p> <p>A value of < 0.5 resets the property to the default.</p> <p>Description: The statistical confidence level for the maximum reducer load factor. Commonly used values other than the default are 0.90 and 0.99.</p>

OraLoader Configuration Properties

OraLoader uses the standard method in Hadoop for specifying configuration properties. These properties can be specified in a configuration file or by using the `-D property=value` option to `GenericOptionsParser` and `ToolRunner`.

[Table 3–4](#) provides brief descriptions of the primary job configuration properties, and [Table 3–5](#) describes the general properties for Oracle Loader for Hadoop. For a complete list and detailed descriptions of all configuration properties, see the `oraloader-conf.xml` document in "[OraLoader for Hadoop Configuration Properties](#)" on page 3-23.

Table 3–4 Primary Job Configuration Properties for Oracle Loader for Hadoop

Property Name	Attributes and Description
<code>oracle.hadoop.loader.jobName</code>	<p>Type: String</p> <p>Default: OraLoader</p> <p>Description: A Hadoop job name for this Oracle loader job. Used as input for the <code>Job.setJobName()</code> method.</p>
<code>oracle.hadoop.loader.targetTable</code>	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: A schema-qualified name for the table to be loaded. Use this option to indicate that all columns of the table are to be loaded and that the names of the input fields match the column names. This property takes precedence over the <code>oracle.hadoop.loader.loaderMapFile</code> property.</p>
<code>oracle.hadoop.loader.loaderMapFile</code>	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Path to the loader map file</p>
<code>oracle.hadoop.loader.tableMetadataFile</code>	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Path to the target table metadata file. Use this option when running in disconnected mode. The table metadata file is created by running the <code>OraLoaderMetadata</code> utility.</p>
<code>oracle.hadoop.loader.olhcachePath</code>	<p>Type: String</p> <p>Default: <code>\${mapred.output.dir}/.../olhcache</code></p> <p>Description: Path to a directory where Oracle Loader for Hadoop can create files that are loaded into the <code>DistributedCache</code>. In distributed mode, the value must be an HDFS path.</p>
<code>oracle.hadoop.loader.extTabDirectoryName</code>	<p>Type: String</p> <p>Default: <code>OLH_EXTTAB_DIR</code></p> <p>Description: The name of the database directory object for the external table's <code>LOCATION</code> data files. This property applies only to the Delimited Text and Data Pump output formats.</p>
<code>oracle.hadoop.loader.sampler.enableSampling</code>	<p>Type: Boolean</p> <p>Default: true</p> <p>Description: Indicates whether the sampling feature is enabled</p>
<code>oracle.hadoop.loader.enableSorting</code>	<p>Type: Boolean</p> <p>Default: true</p> <p>Description: Indicates whether output records within each reducer group should be sorted by the primary key for the table</p>

Table 3–4 (Cont.) Primary Job Configuration Properties for Oracle Loader for Hadoop

Property Name	Attributes and Description
oracle.hadoop.loader.connection.url	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Specifies the URL of the database connection string. This property takes precedence over all other connection properties. If an Oracle wallet is configured as an external password store, then the property value must start with the driver prefix <code>jdbc:oracle:thin:@</code> and the <code>db_connect_string</code> must exactly match the credential defined in the wallet.</p>
oracle.hadoop.loader.connection.user	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Name for database login</p>
oracle.hadoop.loader.connection.password	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Password for the connecting user</p>
oracle.hadoop.loader.connection.wallet_location	<p>Type: String</p> <p>Default: Not defined:</p> <p>Description: File path to an Oracle wallet where the connection information is stored. This property is used only for JDBC connections.</p> <p>For JDBC output format, when an Oracle wallet is used as an external password store, set the following two properties:</p> <ul style="list-style-type: none"> ■ <code>oracle.hadoop.loader.connection.wallet_location</code> ■ <code>oracle.hadoop.loader.connection.url</code> <p>Or, set the following three properties:</p> <ul style="list-style-type: none"> ■ <code>oracle.hadoop.loader.connection.wallet_location</code> ■ <code>oracle.hadoop.loader.connection.tnsEntryName</code> ■ <code>oracle.hadoop.loader.connection.tns_admin</code> <p>For the OCI output format, set the <code>oracle.hadoop.loader.connection.tns_admin</code> property to indicate wallet location.</p> <p>Note that JDBC connections are always made for online loads, even when the OCI Direct Path output format is specified. The same wallet can be used for both connection types.</p>
oracle.hadoop.loader.connection.tnsEntryName	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: Specifies a TNS entry name defined in the <code>tnsnames.ora</code> file. This property is used with the <code>oracle.hadoop.loader.connection.tns_admin</code> property.</p>

Table 3–4 (Cont.) Primary Job Configuration Properties for Oracle Loader for Hadoop

Property Name	Attributes and Description
oracle.hadoop.loader.connection.tns_admin	<p>Type: String</p> <p>Default: Not defined</p> <p>Description: File path to a directory containing Oracle Net configuration files such as sqlnet.ora and tnsnames.ora. If this property is not set, then the value (if any) of the environment variable TNS_ADMIN, is used. Define this property in order to use TNS entry names in database connect strings. This property must be defined when using an Oracle wallet with OCI connections.</p>
oracle.hadoop.loader.connection.defaultExecuteBatch	<p>Type: Integer</p> <p>Default: 100</p> <p>Description: Applicable only for the JDBC and OCI Direct Path output formats. It is the default value for the number of records to be inserted in a batch for each trip to the database. Specify a value greater than 1 to override the default value. If the specified value is less than 1, then this property assumes the default value. Although the maximum value is unlimited, using very large batch sizes is not recommended because it results in a large memory footprint without much increase in performance.</p>
oracle.hadoop.loader.connection.sessionTimeZone	<p>Type: String</p> <p>Default: LOCAL</p> <p>Description: This property is used to alter the session time zone for database connections. Valid values are as follows:</p> <p>[+ -] hh:mm: Hours and minutes before or after UTC</p> <p>Local: The default time zone of the JVM</p> <p>time_zone_region: A valid time zone region</p> <p>This property also determines the default time zone used when parsing input data that is loaded into the following database column types: <code>TIMESTAMP</code>, <code>TIMESTAMP WITH TIME ZONE</code> and <code>TIMESTAMP WITH LOCAL TIME ZONE</code>.</p>
oracle.hadoop.loader.output.dirpathBufsize	<p>Type: Integer</p> <p>Default: 131072</p> <p>Description: This property is used to set the size, in bytes, of the direct path stream buffer for <code>OCIOutputFormat</code>. If needed, values are rounded up to the next nearest multiple of 8 KB.</p>
oracle.hadoop.loader.output.fieldTerminator	<p>Type: String</p> <p>Default: , (comma)</p> <p>Description: A single character to delimit fields for <code>DelimitedTextOutputFormat</code>.</p> <p>Alternate representation: <code>\uHHHHH</code> (where HHHH is the character's UTF-16 encoding).</p>

Table 3–4 (Cont.) Primary Job Configuration Properties for Oracle Loader for Hadoop

Property Name	Attributes and Description
oracle.hadoop.loader.output.initialFieldEncloser	<p>Type: String</p> <p>Default: None</p> <p>Description: When this value is set, fields are always enclosed between the specified character and <code>\${oracle.hadoop.loader.output.trailingFieldEncloser}</code>.</p> <p>If this value is set, it must be either a single character or <code>\uHHHH</code> (where HHHH is the character's UTF-16 encoding). A zero-length value means that there are no enclosers (default value).</p> <p><code>\${oracle.hadoop.loader.output.initialFieldEncloser}</code> and <code>\${oracle.hadoop.loader.output.trailingFieldEncloser}</code> must be either both not set or both set.</p> <p>Use these properties when some field may contain the <code>fieldTerminator</code>. If some field may also contain the <code>trailingFieldEncloser</code>, then set the <code>escapeEnclosers</code> property to true.</p>
oracle.hadoop.loader.output.trailingFieldEncloser	<p>Type: String</p> <p>Default: None</p> <p>Description: When this value is set, fields are always enclosed between <code>\${oracle.hadoop.loader.output.initialFieldEncloser}</code> and the specified character for this property.</p> <p>If this value is set, it must be either a single character, or <code>\uHHHH</code> (where HHHH is the character's UTF-16 encoding). A zero-length value means that there are no enclosers (default value).</p> <p><code>\${oracle.hadoop.loader.output.initialFieldEncloser}</code> and <code>\${oracle.hadoop.loader.output.trailingFieldEncloser}</code> must be either both not set or both set.</p> <p>Use these properties when some field may contain the <code>fieldTerminator</code>. If some field may also contain the <code>trailingFieldEncloser</code>, then set the <code>escapeEnclosers</code> property to true.</p>
oracle.hadoop.loader.output.escapeEnclosers	<p>Type: Boolean</p> <p>Default: false</p> <p>Description: When this is set to true and both initial and trailing field enclosers are set, fields are scanned and embedded trailing enclosure characters are escaped. Use this option when some of the field values may contain the trailing enclosure character.</p>
oracle.hadoop.loader.input.fieldTerminator	<p>Type: String</p> <p>Default: , (comma)</p> <p>Description: A single character to delimit fields for <code>DelimitedTextInputFormat</code>.</p> <p>Alternate representation: <code>\uHHHH</code> (where HHHH is the character's UTF-16 encoding).</p>

Table 3–4 (Cont.) Primary Job Configuration Properties for Oracle Loader for Hadoop

Property Name	Attributes and Description
oracle.hadoop.loader.input.initialFieldEncloser	<p>Type: String</p> <p>Default: None</p> <p>Description: When this value is set, fields can be enclosed between the specified character and <code>\${oracle.hadoop.loader.input.trailingFieldEncloser}</code>.</p> <p>If this value is set, it must be either a single character or <code>\uHHHH</code> (where HHHH is the character's UTF-16 encoding). A zero-length value means no enclosers (default value).</p> <p><code>\${oracle.hadoop.loader.input.initialFieldEncloser}</code> and <code>\${oracle.hadoop.loader.input.trailingFieldEncloser}</code> must be either both not set, or both set.</p>
oracle.hadoop.loader.input.trailingFieldEncloser	<p>Type: String</p> <p>Default: None</p> <p>Description: When this value is set, fields can be enclosed between <code>\${oracle.hadoop.loader.input.initialFieldEncloser}</code> and the specified character.</p> <p>If this value is set, it must be either a single character or <code>\uHHHH</code> (where HHHH is the character's UTF-16 encoding). A zero-length value means no enclosers (default value).</p> <p><code>\${oracle.hadoop.loader.input.initialFieldEncloser}</code> and <code>\${oracle.hadoop.loader.input.trailingFieldEncloser}</code> must be either both not set, or both set.</p>
oracle.hadoop.loader.input.fieldNames	<p>Type: Comma-separated list of strings</p> <p>Default: F0,F1,F2,...</p> <p>Description: Names to assign to input fields. The names are used to create the Avro schema for the record. The strings must be valid JSON name strings.</p>

Table 3–5 General Properties for Oracle Loader for Hadoop

Property Name	Description
<code>mapreduce.inputformat.class</code>	Name of the class implementing <code>InputFormat</code>
<code>mapreduce.outputformat.class</code>	<p>Output options supported by Oracle Loader for Hadoop. The values can be:</p> <ul style="list-style-type: none"> ■ <code>oracle.hadoop.loader.lib.output.DelimitedTextOutputFormat</code> Writes data records to delimited text format files such as comma-separated values (CSV) format files ■ <code>oracle.hadoop.loader.lib.output.JDBCOutputFormat</code> Inserts data records into the target table using JDBC ■ <code>oracle.hadoop.loader.lib.output.OCIOutputFormat</code> Inserts rows into the target table using the Oracle OCI Direct Path interface ■ <code>oracle.hadoop.loader.lib.output.DataPumpOutputFormat</code> Writes rows into binary format files that can be loaded into the target table using an external table

Example of Using Oracle Loader for Hadoop

The example shown in this section uses Oracle Loader for Hadoop in the online database mode using JDBC. It involves the following steps:

1. Create a table in the database. This example uses the `HR.EMPLOYEES` table available as part of the `HR` sample schema in Oracle Database.
2. Implement an `InputFormat` class similar to the examples in the `oracle.hadoop.loader.examples` package.
3. Set the configuration properties. The `MyLoaderMap.xml` document contains the mapping of input data fields to columns in the `HR.EMPLOYEES` table, as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<LOADER_MAP>
<SCHEMA>HR</SCHEMA>
<TABLE>EMPLOYEES</TABLE>
<COLUMN field="empId">EMPLOYEE_ID</COLUMN>
<COLUMN field="lastName">LAST_NAME</COLUMN>
<COLUMN field="email">EMAIL</COLUMN>
<COLUMN field="hireDate" format="MM-dd-yyyy">HIRE_DATE</COLUMN>
<COLUMN field="jobId">JOB_ID</COLUMN>
</LOADER_MAP>
```

The configuration properties in `MyConf.xml` are as follows:

```
<configuration>
  <property>
    <name>mapreduce.inputformat.class</name>
    <value><full_class_name>.MyInputFormat</value>
    <description> Name of the class implementing InputFormat </description>
  </property>

  <property>
    <name>mapreduce.outputformat.class</name>
    <value>oracle.hadoop.loader.lib.output.JDBCOutputFormat</value>
    <description> Output mode after the loader job executes on Hadoop
  </description>
```

```

</property>

<property>
  <name>oracle.hadoop.loader.loaderMapFile</name>
  <value>MyLoaderMap.xml</value>
  <description> The loaderMap file specifying the mapping of input data
    fields to the table columns </description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.user</name>
  <value>HR</value>
  <description> Name of the user connecting to the database</description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.password</name>
  <value>[HR password]</value>
  <description>Password of the user connecting to the database</description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.url</name>
  <value>jdbc:oracle:thin:@//example.com:1521/serviceName</value>
  <description> Database connection string </description>
</property>
</configuration>

```

4. Invoke OraLoader.

```

bin/hadoop jar oraloader.jar oracle.hadoop.loader.OraLoader -libjars \
avro-1.4.1.jar, MyInputFormat.jar -conf MyConf.xml \
-fs [<local|namenode:port>] \
-jt [<local|jobtracker:port>]

```

Target Table Characteristics

Oracle Loader for Hadoop supports loads into a single table, which is referred to as the target table. The target table must exist in the Oracle database. It can contain data or it can be empty.

Supported Data Types

Oracle Loader for Hadoop supports the following database built-in data types:

- VARCHAR2
- CHAR
- NVARCHAR2
- NCHAR
- NUMBER
- FLOAT
- RAW
- BINARY_FLOAT
- BINARY_DOUBLE

- DATE
- TIMESTAMP
- TIMESTAMP WITH TIME ZONE
- TIMESTAMP WITH LOCAL TIME ZONE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND

The target table can contain columns with unsupported data types, but these columns must be nullable, or otherwise set to a value.

Supported Partitioning Strategies

Oracle Loader for Hadoop supports the following single-level partitioning and composite-level partitioning strategies:

- Range
- List
- Hash
- Interval
- Range-Range
- Range-Hash
- Range-List
- List-Range
- List-Hash
- List-List
- Hash-Range
- Hash-Hash
- Hash-List
- Interval-Range
- Interval-Hash
- Interval-List

Oracle Loader for Hadoop does not support reference partitioning or virtual column-based partitioning.

Loader Map XML Schema Definition

This section contains the XML schema definition (XSD) for the loader map that specifies the columns to be loaded into the target table:

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema elementFormDefault="qualified" attributeFormDefault="unqualified"
  xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:attributeGroup name="columnAttrs">
    <xs:annotation>
      <xs:documentation>Column attributes define how to map input fields to the
        database column. field - is the name of the field in the
        IndexedRecord input object. The field name need not be
```

```

        unique. This means that the same input field can map to
        different columns in the database table. format - is a
        format string for interpreting the input. For example,
        if the field is a date then the format is a date format
        string suitable for interpreting dates</xs:documentation>
    </xs:annotation>
    <xs:attribute name="field" type="xs:token" use="optional"/>
    <xs:attribute name="format" type="xs:token" use="optional"/>
</xs:attributeGroup>
<xs:simpleType name="TOKEN_T">
    <xs:restriction base="xs:token">
        <xs:minLength value="1"/>
    </xs:restriction>
</xs:simpleType>
<xs:element name="LOADER_MAP">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="SCHEMA" type="TOKEN_T" minOccurs="0"/>
            <xs:element name="TABLE" type="TOKEN_T" nillable="false"/>
            <xs:element name="COLUMN" maxOccurs="unbounded" minOccurs="0">
                <xs:annotation>
                    <xs:documentation>specifies the database column name that will be
                    loaded. Each column name must be unique.
                </xs:documentation>
            </xs:annotation>
            <xs:complexType>
                <xs:simpleContent>
                    <xs:extension base="TOKEN_T">
                        <xs:attributeGroup ref="columnAttrs"/>
                    </xs:extension>
                </xs:simpleContent>
            </xs:complexType>
        </xs:element>
    </xs:sequence>
</xs:complexType>
</xs:element>
</xs:schema>

```

OraLoader for Hadoop Configuration Properties

This is the `oraloader-conf.xml` document, which describes the configuration properties for Oracle Loader for Hadoop:

```

<?xml version="1.0"?>
<!--
Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.

NAME
    oraloader-conf.xml

DESCRIPTION
    Config properties for OLH.

    This file is loaded as the very first conf resource.
    Properties without default values are commented out.
-->
<configuration>
    <property>
        <name>oracle.hadoop.loader.libjars</name>
        <value>${oracle.hadoop.loader.olh_home}/jlib/ojdbc6.jar,

```

```

${oracle.hadoop.loader.olh_home}/jlib/orai18n.jar,
${oracle.hadoop.loader.olh_home}/jlib/orai18n-utility.jar,
${oracle.hadoop.loader.olh_home}/jlib/orai18n-mapping.jar,
${oracle.hadoop.loader.olh_home}/jlib/orai18n-collation.jar,
${oracle.hadoop.loader.olh_home}/jlib/oraclepki.jar,
${oracle.hadoop.loader.olh_home}/jlib/osdt_cert.jar,
${oracle.hadoop.loader.olh_home}/jlib/osdt_core.jar,
${oracle.hadoop.loader.olh_home}/jlib/commons-math-2.2.jar,
${oracle.hadoop.loader.olh_home}/jlib/jackson-core-asl-1.5.2.jar,
${oracle.hadoop.loader.olh_home}/jlib/jackson-mapper-asl-1.5.2.jar,
${oracle.hadoop.loader.olh_home}/jlib/avro-1.5.4.jar,
${oracle.hadoop.loader.olh_home}/jlib/avro-mapred-1.5.4.jar</value>
  <description>Comma separated list of library jar files. These jars get
    injected into the command-line arguments under the
    GenericOptionsParser's "-libjars" option. When a "-libjars"
    option is used as a command-line argument, then this list of
    jars is prepended to the list following "-libjars". Users can
    distribute their application jars using this property in place
    of, or in combination with, the "-libjars" option.</description>
</property>

<property>
  <name>oracle.hadoop.loader.sharedLibs</name>
  <value>${oracle.hadoop.loader.olh_home}/lib/libolh11.so,
${oracle.hadoop.loader.olh_home}/lib/libclntsh.so.11.1,
${oracle.hadoop.loader.olh_home}/lib/libnnz11.so,
${oracle.hadoop.loader.olh_home}/lib/libociei.so</value>
</property>

<property>
  <name>oracle.hadoop.loader.olh_home</name>
  <value/>
  <description>
    A path to the OLH_HOME on the node where the OraLoader job
    is initiated. OraLoader uses this path to locate required libraries.
    If this property is not set, OraLoader will use the value in the environment
    variable OLH_HOME.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.jobName</name>
  <value>OraLoader</value>
  <description>
    Hadoop job name for this Oracle loader job. Used as input for
    the Job.setJobName() method.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.targetTable</name>
  <value/>
  <description>
    A schema qualified name for the table to be loaded. Use this
    property to indicate that all columns of the table will be
    loaded and that the names of the input fields match the
    column names. This property takes precedence over the
    oracle.hadoop.loader.loaderMapFile property. The default
    value is null.
  </description>

```



```
</property>

<property>
  <name>oracle.hadoop.loader.loaderMapFile</name>
  <value/>
  <description>
    Path to the loader map file. Use a file:// schema to indicate a local file.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.tableMetadataFile</name>
  <value/>
  <description>
    Path to the target table metadata file. Use this property when
    running in disconnected mode. The table metadata file is
    created by running the OraLoaderMetadata utility.
    Use a file:// schema to indicate a local file.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.olhcachePath</name>
  <value>${mapred.output.dir}/../olhcache</value>
  <description>
    Path to a directory where Oracle Loader for Hadoop can create
    files that will be loaded into the DistributedCache.
    Unique file names are generated every time; one may want to empty it,
    or it will grow bigger and bigger if jobs are run
    using the same olhcache directory.

    The default value is a directory called 'olhcache' in the parent directory
    of the job's output directory (i.e. ${mapred.output.dir}).

    In distributed mode, the value must be a hdfs path
    (see javaDoc for org.apache.hadoop.filecache.DistributedCache).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.loadByPartition</name>
  <value>true</value>
  <description>
    Instructs the output format to perform a partition-aware load.
    For DelimitedText output format, this option controls whether the
    keyword "PARTITION" appears in the generated .ctl file(s).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.extTabDirectoryName</name>
  <value>OLH_EXTTAB_DIR</value>
  <description>
    The name of the Oracle directory object for the external table's
    LOCATION data files. This property applies only to the DelimitedText
    and DataPump output formats.
  </description>
</property>

<property>
```

```

<name>oracle.hadoop.loader.sampler.enableSampling</name>
<value>>true</value>
<description>
  Indicates whether the sampling feature is enabled.
  Set the value to false to disable this feature.
</description>
</property>

<property>
  <name>oracle.hadoop.loader.enableSorting</name>
  <value>true</value>
  <description>
    Indicates whether output records within each reducer group
    should be sorted by the primary key for the table.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.configuredCounters</name>
  <value>MAPPER,OUTPUT</value>
  <description>
    Turns ON Oracle Loader for Hadoop counters by category. The value is a
    comma separated list of zero or more of the following keywords:
    MAPPER, REDUCER, OUTPUT, and SAMPLER.

    Note that the input error counters (displayed in the
    "map phase counters" section of the final report) are always on,
    regardless of the presence of the keyword MAPPER in this list.

    Newer release of Hadoop (0.20.203, cdh3u3) impose a hard limit on the
    total number of counters a job can use (see property
    mapreduce.job.counters.limit in mapred-site.xml). Note that this
    limit cannot be changed on a per-job basis, and the cluster needs
    to be restarted after the property has been updated on all nodes.

    In order to turn off all the Oracle Loader for Hadoop specific counters,
    set this property's value to an empty list using either:

    -D oracle.hadoop.loader.configuredCounters=

    or

    <property>
      <name>oracle.hadoop.loader.configuredCounters</name>
      <value>,</value>
    </property>

  </description>
</property>

<!-- CONNECTION properties -->

<property>
  <name>oracle.hadoop.loader.connection.url</name>
  <value/>
  <description>
    Specifies the URL of the database connection string. This property
    takes precedence and overrides all other connection properties.

    If Oracle Wallet is configured as an external password store,

```

the property value must start with the driver prefix: jdbc:oracle:thin:@ and the db_connect_string must exactly match the credential defined in the wallet.

Example 1: (using oracle net syntax)

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=
  (ADDRESS=(PROTOCOL=TCP) (HOST=myhost) (PORT=1521)))
  (CONNECT_DATA=(SERVICE_NAME=my_db_service_name)))
```

Example 2: (using TNS entry)

```
jdbc:oracle:thin:@myTNS
```

- Also see documentation for
oracle.hadoop.loader.connection.wallet_location

If Oracle Wallet is NOT used, then set the following conf properties:
oracle.hadoop.loader.connection.url

Examples of connection URL styles:

thin-style:

```
jdbc:oracle:thin:@//myhost:1521/my_db_service_name
jdbc:oracle:thin:user/password@//myhost:1521/my_db_service_name
```

Oracle Net:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=
  (ADDRESS=(PROTOCOL=TCP) (HOST=myhost) (PORT=1521)))
  (CONNECT_DATA=(SERVICE_NAME=my_db_service_name)))
```

TNSEntry Name:

```
jdbc:oracle:thin:@myTNSEntryName
```

AND

```
oracle.hadoop.loader.connection.user
oracle.hadoop.loader.connection.password
```

If OCIOutputFormat is configured, and Oracle Wallet is not used, then username and password must be specified in these separate properties.

```
</description>
</property>
```

```
<property>
  <name>oracle.hadoop.loader.connection.user</name>
  <value/>
  <description>Name for the database login.</description>
</property>
```

```
<property>
  <name>oracle.hadoop.loader.connection.password</name>
  <value/>
  <description>Password for the connecting user.</description>
</property>
```

```
<property>
  <name>oracle.hadoop.loader.connection.wallet_location</name>
  <value/>
  <description>File path to an Oracle wallet where the connection information
  is stored. This property is used only for JDBC connections. For JDBC output
  format, when using Oracle Wallet as an external password store, set the
  following two properties:
```

- oracle.hadoop.loader.connection.wallet_location
- oracle.hadoop.loader.connection.url

Or, set the following three properties:

- oracle.hadoop.loader.connection.wallet_location
- oracle.hadoop.loader.connection.tnsEntryName
- oracle.hadoop.loader.connection.tns_admin

For the OCI output format, set the oracle.hadoop.loader.connection.tns_admin property to indicate wallet location.

Note that JDBC connections are always made for online loads, even when the OCI Direct Path output format is specified. The same wallet can be used for both connection types.

```

</description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.tnsEntryName</name>
  <value/>
  <description>Specifies a TNS entry name defined in the tnsnames.ora file.
  This property is used together with the
  oracle.hadoop.loader.connection.tns_admin property.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.tns_admin</name>
  <value/>
  <description>File path to a directory containing
  SQL*Net configuration files like sqlnet.ora and tnsnames.ora.
  If this property is not set, the value of the environment
  variable TNS_ADMIN will be used. Define this property in order
  to use TNS entry names in database connect strings.
  This property must be defined when using an Oracle Wallet with OCI
  connections.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.defaultExecuteBatch</name>
  <value>100</value>
  <description>
    Applicable only for JDBC and OCI output formats. The default
    value for the number of records to be inserted in a batch for
    each trip to the database. Specify a value >= 1 to
    override the default value. If the specified value is less than 1,
    this property assumes the default value. Though the maximum
    value is unlimited, using very large batch sizes is not
    recommended, as it results in a large memory footprint without
    much increase in performance.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.connection.sessionTimezone</name>
  <value>LOCAL</value>

```

```

<description>
  This property is used to alter the session time zone for
  database connections. Valid values are:

      [+|-] hh:mm      - hours and minutes before or after UTC
      LOCAL            - the default timezone of the JVM
      time_zone_region - a valid time zone region

  This property also determines the default timezone when parsing
  input data that will be loaded to database column types:
  TIMESTAMP, TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE
</description>
</property>

<!-- properties for OCIOutputFormat -->
<property>
  <name>oracle.hadoop.loader.output.dirpathBufsize</name>
  <value>131072</value>
  <description>
    This property is used to set the size, in bytes, of the direct path stream
    buffer for OCIOutputFormat. If needed, values are rounded up to the next
    nearest multiple of 8k.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.compressionFactors</name>
  <value>BASIC=5.0,OLTP=5.0,QUERY_LOW=10.0,QUERY_HIGH=10.0,
    ARCHIVE_LOW=10.0,ARCHIVE_HIGH=10.0</value>
  <description>
    This property is used to define the compression factor for different types
    of compression. The format is a comma separated list of name=value pairs
    where name is one of BASIC, OLTP, QUERY_LOW, QUERY_HIGH, ARCHIVE_LOW, or
    ARCHIVE_HIGH. Value is a decimal number.
  </description>
</property>

<!-- properties for DelimitedTextOutputFormat -->
<property>
  <name>oracle.hadoop.loader.output.fieldTerminator</name>
  <value></value>
  <description>
    A single character to delimit fields for DelimitedTextOutputFormat.
    Alternate representation: \uHHHH (where HHHH is the character's UTF-16
    encoding).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.output.initialFieldEncloser</name>
  <value></value>
  <description>
    When this value is set, fields are always enclosed between the
    specified character and
    ${oracle.hadoop.loader.output.trailingFieldEncloser}.

    If this value is set, it must be either a single character, or \uHHHH
    (where HHHH is the character's UTF-16 encoding).

    ${oracle.hadoop.loader.output.initialFieldEncloser} and

```

```
    ${oracle.hadoop.loader.output.trailingFieldEncloser} must be either
    both not set, or both set.
    A zero length value means no enclosers (default value).

    Use this when some field may contain the fieldTerminator.
    If some field may also contain the trailingFieldEncloser, then
    the escapeEnclosers property should be set to true.
</description>
</property>

<property>
  <name>oracle.hadoop.loader.output.trailingFieldEncloser</name>
  <value></value>
  <description>
    When this value is set, fields are always enclosed between
    ${oracle.hadoop.loader.output.initialFieldEncloser} and the
    specified character for this property.

    If this value is set, it must be either a single character, or \uHHHH
    (where HHHH is the character's UTF-16 encoding).

    ${oracle.hadoop.loader.output.initialFieldEncloser} and
    ${oracle.hadoop.loader.output.trailingFieldEncloser} must be either
    both not set, or both set.
    A zero length value means no enclosers (default value).

    Use this when some field may contain the fieldTerminator.
    If some field may also contain the trailingFieldEncloser, then
    the escapeEnclosers property should be set to true.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.output.escapeEnclosers</name>
  <value>>false</value>
  <description>
    When this is set to true and both initial and trailing field enclosers
    are set, fields will be scanned, and embedded trailing encloser
    characters will be escaped. Use this option when some of the field
    values may contain the trailing encloser character.
  </description>
</property>

<!-- properties for DelimitedTextInputFormat -->
<property>
  <name>oracle.hadoop.loader.input.fieldNames</name>
  <value/>
  <description>
    Comma-separated list of names to assign to input fields.
    The names are used to create the Avro schema for the record.
    The strings must be valid JSON name strings.

    If this property is not set, the names F0,F1,F2,... will be used
    (consistent with oracle.hadoop.loader.examples.CSVInputFormat).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.input.fieldTerminator</name>
  <value>,</value>
```

```

<description>
  A single character to delimit fields for DelimitedTextInputFormat.
  Alternate representation: \uHHHH (where HHHH is the character's UTF-16
  encoding).
</description>
</property>

<property>
  <name>oracle.hadoop.loader.input.initialFieldEncloser</name>
  <value></value>
  <description>
    When this value is set, fields are allowed to be enclosed
    between the specified character and
    ${oracle.hadoop.loader.input.trailingFieldEncloser}.

    If this value is set, it must be either a single character, or \uHHHH
    (where HHHH is the character's UTF-16 encoding).

    ${oracle.hadoop.loader.input.initialFieldEncloser} and
    ${oracle.hadoop.loader.input.trailingFieldEncloser} must be either
    both not set, or both set.
    A zero length value means no enclosers (default value).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.input.trailingFieldEncloser</name>
  <value></value>
  <description>
    When this value is set, fields are allowed to be enclosed
    between ${oracle.hadoop.loader.input.initialFieldEncloser}
    and the specified character.

    If this value is set, it must be either a single character, or \uHHHH
    (where HHHH is the character's UTF-16 encoding).

    ${oracle.hadoop.loader.input.initialFieldEncloser} and
    ${oracle.hadoop.loader.input.trailingFieldEncloser} must be either
    both not set, or both set.
    A zero length value means no enclosers (default value).
  </description>
</property>

<!--Properties for tuning the sampler-->
<!-- set numThreads > 1 for large datasets -->
<property>
  <name>oracle.hadoop.loader.sampler.numThreads</name>
  <value>5</value>
  <description>Number of sampler threads.
  This value should be set based on the processor and memory resources
  available to the job tracker node. A higher number of sampler threads
  implies higher concurrency in sampling.
  The default value is 5 threads.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.maxLoadFactor</name>
  <value>0.05</value>
  <description>

```

The maximum acceptable reducer load factor.
In a perfectly load balanced job, every reducer is assigned an equal amount of work (or load).
Load factor is the percent overload per reducer i.e. (assigned load - ideal load)%
For example: a value of 0.05, indicates that it is acceptable for reducers to be assigned up to 5% more data than their ideal load.
If load balancing is successful, it guarantees this maximum load factor at the specified confidence.
(see oracle.hadoop.loader.sampler.loadCI)
Default = 0.05, another common value is 0.1.

```
</description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.loadCI</name>
  <value>0.95</value>
  <description>
    The confidence level for the specified
    maximum reducer load factor.
    (See oracle.hadoop.loader.sampler.maxLoadFactor)
    Default = 0.95, other common values = 0.90, 0.99
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.minSplits</name>
  <value>5</value>
  <description>
    The minimum number of splits that will be
    read by the sampler. If the total number of splits
    is lesser than this value, then the sampler will read
    all splits. Splits may be read partially.
    A non-positive value is equivalent to minSplits=1.
    The default value is 5.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.hintMaxSplitSize</name>
  <value>1048576</value>
  <description>
    The sampler sets Hadoop configuration parameter
    mapred.max.split.size to this value before it calls the InputFormat's
    getSplits() method.
    The value of mapred.max.split.size is only set to this value for the
    duration of sampling, it is not changed in the actual job
    configuration. Some InputFormats (e.g. FileInputFormat) use the
    maximum split size as a hint to determine the number of splits
    returned by getSplits(). Smaller split sizes imply that more
    chunks of data will be sampled at random (good). While large splits are
    better for IO performance, they are not necessarily better for sampling.
    Set this value to be small enough for good sampling performance,
    but not any smaller: extremely small values can cause inefficient IO
    performance and cause getSplits() to run out of memory by returning too
    many splits.
    The recommended minimum value for this property is 1048576 bytes (1 MB).
    This value can be increased for larger datasets (e.g. tens of terabytes)
    or if the InputFormat's getSplits() method throws an OutOfMemoryError.
    If the specified value is less than 1, this property is ignored.
```



```

    The default value is 1048576 bytes (1 MB).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.hintNumMapTasks</name>
  <value>100</value>
  <description>
    The sampler sets Hadoop configuration parameter
    mapred.map.tasks to this value for the duration of sampling.
    The value of mapred.map.tasks is not changed in the actual job
    configuration. Some InputFormats (e.g. DBInputFormat) use the
    number of map tasks parameter as a hint to determine the number of
    splits returned by getSplits(). Higher values imply that more chunks
    of data will be sampled at random (good). The default value is 100.
    This value should typically be increased for large datasets (e.g. more
    than a million rows), while keeping in mind that extremely large values
    can cause the InputFormat's getSplits() method to run out of memory by
    returning too many splits.
    If the specified value is less than 1, this property is ignored.
    The default value is 100.
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.maxSamplesPct</name>
  <value>0.01</value>
  <description>
    This property specifies the maximum data to sample, as a
    percentage of the total amount of data. In general, the
    sampler will stop sampling if any one of the following is true:
    (1) it has collected the minimum number of samples
        required for optimal load-balancing, or
    (2) the percent of data sampled exceeds
        oracle.hadoop.loader.sampler.maxSamplesPct, or
    (3) the number of bytes sampled exceeds
        oracle.hadoop.loader.sampler.maxHeapBytes.
    If this parameter is set to a negative value,
    condition (2) is not imposed.
    The default value is 0.01 (1%).
  </description>
</property>

<property>
  <name>oracle.hadoop.loader.sampler.maxHeapBytes</name>
  <value>-1</value>
  <description>
    This value specifies the maximum memory available to
    the sampler in bytes. In general, the sampler will
    stop sampling when any one of these conditions is true:
    (1) it has collected the minimum number of samples
        required for optimal load-balancing, or
    (2) the percent of data sampled exceeds
        oracle.hadoop.loader.sampler.maxSamplesPct, or
    (3) the number of bytes sampled exceeds
        oracle.hadoop.loader.sampler.maxHeapBytes.
    If this parameter is set to a negative value,
    condition (3) is not imposed.
    Default = -1 (no memory restrictions on the sampler).
  </description>

```

```
</property>  
</configuration>
```

Third-Party Licenses for Bundled Software

Oracle Loader for Hadoop installs the following third-party products:

- Apache Avro
- Apache Commons Mathematics Library
- Jackson JSON Processor

Oracle Loader for Hadoop includes Oracle 11g Release 2 (11.2) client libraries. For information about third party product included with Oracle Database 11g Release 2 (11.2), refer to *Oracle Database Licensing Information*.

Unless otherwise specifically noted, or as required under the terms of the third party license (e.g., LGPL), the licenses and statements herein, including all statements regarding Apache-licensed code, are intended as notices only.

Apache Licensed Code

The following is included as a notice in compliance with the terms of the Apache 2.0 License, and applies to all programs licensed under the Apache 2.0 license:

You may not use the identified files except in compliance with the Apache License, Version 2.0 (the "License.")

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

A copy of the license is also reproduced below.

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or

otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent

infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
 - c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Do not include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>) (listed below):

Apache Avro avro-1.5.4.jar

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache Commons Mathematics Library 2.2

Copyright 2001-2011 The Apache Software Foundation

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Jackon JSON Library 1.5.2

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Oracle Data Integrator Application Adapter for Hadoop

This chapter describes how to use the knowledge modules in Oracle Data Integrator (ODI) Application Adapter for Hadoop. It contains the following sections:

- [Introduction](#)
- [Setting Up the Topology](#)
- [Setting Up an Integration Project](#)
- [Creating an Oracle Data Integrator Model from a Reverse-Engineered Hive Model](#)
- [Designing the Interface](#)

See Also: *Oracle Fusion Middleware Application Adapters Guide for Oracle Data Integrator*

Introduction

Apache Hadoop is designed to handle and process data that is typically from data sources that are nonrelational and data volumes that are beyond what is handled by relational databases.

Oracle Data Integrator (ODI) Application Adapter for Hadoop enables data integration developers to integrate and transform data easily within Hadoop using Oracle Data Integrator. Employing familiar and easy-to-use tools and preconfigured knowledge modules (KMs), the application adapter provides the following capabilities:

- Loading data into Hadoop from the local file system and HDFS
- Performing validation and transformation of data within Hadoop
- Loading processed data from Hadoop to an Oracle database for further processing and generating reports

Knowledge modules (KMs) contain the information needed by Oracle Data Integrator to perform a specific set of tasks against a specific technology. An application adapter is a group of knowledge modules. Thus, ODI Application Adapter for Hadoop is a group of knowledge modules for accessing data stored in Hadoop.

Concepts

Typical processing in Hadoop includes data validation and transformations that are programmed as MapReduce jobs. Designing and implementing a MapReduce job requires expert programming knowledge. However, when you use Oracle Data

Integrator and ODI Application Adapter for Hadoop, you do not need to write MapReduce jobs. Oracle Data Integrator uses Hive and the Hive Query Language (HiveQL), a SQL-like language for implementing MapReduce jobs.

When you implement a big data processing scenario, the first step is to load the data into Hadoop. The data source is typically in the local file system, HDFS, Hive tables, or external Hive tables.

After the data is loaded, you can validate and transform it by using HiveQL like you use SQL. You can perform data validation (such as checking for NULLS and primary keys), and transformations (such as filtering, aggregations, set operations, and derived tables). You can also include customized procedural snippets (scripts) for processing the data.

When the data has been aggregated, condensed, or processed into a smaller data set, you can load it into an Oracle database for further processing and analysis. Oracle Loader for Hadoop is recommended for optimal loading into an Oracle database.

Knowledge Modules

Oracle Data Integrator provides the knowledge modules (KMs) described in [Table 4-1](#) for use with Hadoop.

Table 4-1 ODI Application Adapter for Hadoop Knowledge Modules

KM Name	Description	Source	Target
IKM File to Hive (Load Data)	Loads data from local and HDFS files into Hive tables. It provides options for better performance through Hive partitioning and fewer data movements. This knowledge module supports wildcards (*,?).	File system	Hive
IKM Hive Control Append	Integrates data into a Hive target table in truncate/insert (append) mode. Data can be controlled (validated). Invalid data is isolated in an error table and can be recycled.	Hive	Hive
IKM Hive Transform	Integrates data into a Hive target table after the data has been transformed by a customized script such as Perl or Python	Hive	Hive
IKM File-Hive to Oracle (OLH)	Integrates data from an HDFS file or Hive source into an Oracle database target using Oracle Loader for Hadoop	File system or Hive	Oracle Database
CKM Hive	Validates data against constraints	NA	Hive
RKM Hive	Reverse engineers Hive tables	Hive metadata	NA

Security

For security information for Oracle Data Integrator, see the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.

Setting Up the Topology

To set up the topology in Oracle Data Integrator, you identify the data server and the physical and logical schemas that are used to store the file system and Hive information.

This section contains the following topics:

- [Setting Up File Data Sources](#)
- [Setting Up Hive Data Sources](#)
- [Setting Up the Oracle Data Integrator Agent to Execute Hadoop Jobs](#)
- [Configuring Oracle Data Integrator Studio for Executing Hadoop Jobs on the Local Agent](#)

Setting Up File Data Sources

In the Hadoop context, there is a distinction between files in Hadoop Distributed File System (HDFS) and local files (files outside of HDFS).

To define a data source:

1. Create a DataServer object under File technology.
2. Create a Physical Schema object for every directory to be accessed.
3. Create a Logical Schema object for every directory to be accessed.
4. Create a Model for every Logical Schema.
5. Create one or more data stores for each different type of file and wildcard name pattern.
6. For HDFS files, create a DataServer object under File technology by entering the HDFS name node in the field JDBC URL. For example:

```
hdfs://bda1node01.example.com:8020
```

Note: No dedicated technology is defined for HDFS files.

Setting Up Hive Data Sources

The following steps in Oracle Data Integrator are required for connecting to a Hive system. Oracle Data Integrator connects to Hive by using JDBC.

Prerequisites

The Hive technology must be included in the standard Oracle Data Integrator technologies. If it is not, then import the technology in `INSERT_UPDATE` mode from the xml-reference directory.

You must add all Hive-specific flex fields. For pre-11.1.1.6.0 repositories, the flex fields are added during the repository upgrade process.

To set up a Hive data source:

1. Place all required Hive JDBC JAR files into the Oracle Data Integrator user lib folder:

```
$HIVE_HOME/lib/*.jar
$HADOOP_HOME/hadoop-*-core*.jar,
```

```
$HADOOP_HOME/Hadoop-*-tools*.jar
```

2. Create a DataServer object under Hive technology.
3. Set the following locations under JDBC:
 - JDBC Driver: `org.apache.hadoop.hive.jdbc.HiveDriver`
 - JDBC URL: for example, `jdbc:hive://BDA:10000/default`
4. Set the following under Flexfields:
 - Hive Metastore URIs: for example, `thrift://BDA:10000`
5. Create a Physical Default Schema.
 - As of Hive 0.7.0, no schemas or databases are supported. Only Default is supported. Enter `default` in both schema fields of the physical schema definition.
6. Ensure that the Hive server is up and running.
7. Test the connection to the DataServer.
8. Create a Logical Schema object.
9. Create at least one Model for the LogicalSchema.
10. Import RKM Hive as a global knowledge module or into a project.
11. Create a new model for Hive Technology pointing to the logical schema.
12. Perform a custom reverse-engineering operation using RKM Hive.

At the end of this process, the Hive DataModel contains all Hive tables with their columns, partitioning, and clustering details stored as flex field values.

Setting Up the Oracle Data Integrator Agent to Execute Hadoop Jobs

After setting up an Oracle Data Integrator agent, configure it to work with ODI Application Adapter for Hadoop.

To configure the Oracle Data Integrator agent:

1. Install Hadoop on your Oracle Data Integrator agent computer. Ensure that the `HADOOP_HOME` environment variable is set.
 - For Oracle Big Data Appliance, see the *Oracle Big Data Appliance Software User's Guide* for instructions for setting up a remote Hadoop client.
2. Install Hive on your Oracle Data Integrator agent computer. Ensure that the `HIVE_HOME` environment variable is set.
3. Copy these JAR files to the Oracle Data Integrator agent drivers directory.

```
$HIVE_HOME/lib/*.jar,  
$HADOOP_HOME/hadoop-*-core*.jar  
$HADOOP_HOME/hadoop-*-tools*.jar
```

See the *Oracle Fusion Middleware Installation Guide for Oracle Data Integrator* for information about adding drivers to the agent. This step enables the Oracle Data Integrator agent to load the Hive JDBC driver.

4. Set environment variable `ODI_HIVE_SESSION_JARS` to include Hive Regex SerDe:
 - `ODI_HIVE_SESSION_JARS=$HIVE_HOME/lib/hive-contrib-0.7.1-cdh3u3.jar`

Include other JAR files as required, such as custom SerDes JAR files. These JAR files are added to every Hive JDBC session and thus are added to every Hive MapReduce job.

5. Set environment variable HADOOP_CLASSPATH:

```
HADOOP_CLASSPATH=$HIVE_HOME/lib/hive-metastore-0.7.1-cdh3u3.jar:$HIVE_
HOME/lib/libthrift.jar:$HIVE_HOME/lib/libfb303.jar:$HIVE_
HOME/lib/hive-common-0.7.1-cdh3u3.jar:$HIVE_
HOME/lib/hive-exec-0.7.1-cdh3u3.jar.
```

This setting enables the Hadoop script to start Hive MapReduce jobs.

To use Oracle Loader for Hadoop:

1. Install Oracle Loader for Hadoop on your Oracle Data Integrator agent system. See ["Installing Oracle Loader for Hadoop"](#) on page 1-6.
2. Install Oracle Database client on your Oracle Data Integrator agent system. See ["Oracle Loader for Hadoop Setup"](#) on page 1-6 for the Oracle Database client version.
3. Set environment variable OLH_HOME.
4. Set environment variable ODI_OLH_JARS.

You must list all JAR files required for Oracle Loader for Hadoop. See the `oracle.hadoop.loader.libjars` property in ["OraLoader for Hadoop Configuration Properties"](#) on page 3-23.

This is a comma-separated list of JAR files:

```
ODI_OLH_JARS=OLH_HOME/jlib/ojdbc6.jar,$OLH_HOME/jlib/orai18n.jar,$OLH_
HOME/jlib/orai18n-utility.jar,$OLH_HOME/jlib/orai18n-mapping.jar,$OLH_
HOME/jlib/orai18n-collation.jar,$OLH_HOME/jlib/oraclepki.jar,$OLH_
HOME/jlib/osdt_cert.jar,$OLH_HOME/jlib/osdt_core.jar,$OLH_
HOME/jlib/commons-math-2.2.jar,$OLH_HOME/jlib/jackson-core-asl-1.5.2.jar,$OLH_
HOME/jlib/jackson-mapper-asl-1.5.2.jar,$OLH_HOME/jlib/avro-1.5.4.jar,$OLH_
HOME/jlib/avro-mapred-1.5.4.jar,$OLH_HOME/jlib/oraloder.jar,$HIVE_
HOME/lib/hive-metastore-0.7.1-cdh3u3.jar,$HIVE_HOME/lib/libthrift.jar,$HIVE_
HOME/lib/libfb303.jar,$HIVE_HOME/lib/hive-common-0.7.1-cdh3u3.jar,$HIVE_
HOME/lib/hive-exec-0.7.1-cdh3u3.jar
```

5. Add paths to HADOOP_CLASSPATH:

```
$HADOOP_CLASSPATH= $OLH_HOME/jlib:$HADOOP_CLASSPATH
```

6. Verify that ODI_OLH_SHAREDLIBS lists all native libraries required for Oracle Loader for Hadoop. See the `oracle.hadoop.loader.sharedLibs` property in the ["OraLoader for Hadoop Configuration Properties"](#) on page 3-23.

This is a comma-separated list of shared libraries files:

```
ODI_OLH_SHAREDLIBS= $OLH_HOME/lib/libolh11.so,$OLH_
HOME/lib/libclntsh.so.11.1,$OLH_HOME/lib/libnzn11.so,$OLH_HOME/lib/libociei.so
```

7. If Oracle Loader for Hadoop is used in OCI Direct Path output format, then check these variables:
 - JAVA_LIBRARY_PATH must include \$OLH_HOME/lib. It locates the Oracle Loader for Hadoop native library libolh11.so.
 - LD_LIBRARY_PATH must include \$ORACLE_HOME/lib.

Configuring Oracle Data Integrator Studio for Executing Hadoop Jobs on the Local Agent

For executing Hadoop jobs on the local agent of an Oracle Data Integrator Studio installation, follow the configuration steps in the previous section with the following change: Copy JAR files into the Oracle Data Integrator userlib directory instead of the drivers directory.

Setting Up an Integration Project

Setting up a project follows the standard procedures. See the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.

Import the following KMs into Oracle Data Integrator project:

- IKM File to Hive (Load Data)
- IKM Hive Control Append
- IKM Hive Transform
- IKM File-Hive to Oracle (OLH)
- CKM Hive
- RKM Hive

Creating an Oracle Data Integrator Model from a Reverse-Engineered Hive Model

This section contains the following topics:

- [Creating a Model](#)
- [Reverse Engineering Hive Tables](#)

Creating a Model

To create a model that is based on the technology hosting Hive and on the logical schema created when you configured the Hive connection, follow the standard procedure described in the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.

Reverse Engineering Hive Tables

RKM Hive is used to reverse engineer Hive tables and views. To perform a customized reverse-engineering of Hive tables with RKM Hive, follow the usual procedures, as described in the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*. This topic details information specific to Hive tables.

The reverse-engineering process creates the data stores for the corresponding Hive table or views. You can use the data stores as either a source or a target in an integration interface.

RKM Hive

RKM Hive reverses these metadata elements:

- Hive tables and views as Oracle Data Integrator data stores.

Specify the reverse mask in the Mask field, and then select the tables and views to reverse. The Mask field in the Reverse tab filters reverse-engineered objects based on their names. The Mask field cannot be empty and must contain at least the percent sign (%).

- Hive columns as Oracle Data Integrator columns with their data types.
- Information about buckets, partitioning, clusters, and sort columns are set in the respective flex fields in the Oracle Data Integrator data store or column metadata.

Table 4–2 describes the options for RKM Hive.

Table 4–2 RKM Hive Options

Option	Description
USE_LOG	Log intermediate results?
LOG_FILE_NAME	Path and file name of log file. Default path is the user home and the default file name is reverse.log.

Table 4–3 describes the created flex fields.

Table 4–3 Flex Fields for Reverse-Engineered Hive Tables and Views

Object	Flex Field Name	Flex Field Code	Flex Field Type	Description
DataStore	Hive Buckets	HIVE_BUCKETS	String	Number of buckets to be used for clustering
Column	Hive Partition Column	HIVE_PARTITION_COLUMN	Numeric	All partitioning columns are marked as "1". Partition information can come from the following: <ul style="list-style-type: none"> ■ Mapped source column ■ Constant value specified in the target column ■ File name fragment
Column	Hive Cluster Column	HIVE_CLUSTER_COLUMN	Numeric	All cluster columns are marked as "1".
Column	Hive Sort Column	HIVE_SORT_COLUMN	Numeric	All sort columns are marked as "1".

Designing the Interface

After reverse engineering Hive tables and configuring them, you can choose from these interface configurations:

- [Loading Data from Files into Hive](#)
- [Validating and Transforming Data Within Hive](#)
- [Loading Data into an Oracle Database from Hive and HDFS](#)

Loading Data from Files into Hive

To load data from the local file system or the HDFS file system into Hive tables:

1. Create the data stores for local files and HDFS files.

Refer to the *Oracle Fusion Middleware Connectivity and Knowledge Modules Guide for Oracle Data Integrator* for information about reverse engineering and configuring local file data sources.

2. Create an interface using the file data store as the source and the corresponding Hive table as the target. Use the IKM File to Hive (Load Data) knowledge module specified in the flow tab of the interface. This integration knowledge module loads data from flat files into Hive, replacing or appending any existing data.

IKM File to Hive

IKM File to Hive (Load Data) supports:

- One or more input files. To load multiple source files, enter an asterisk or a question mark as a wildcard character in the resource name of the file DataStore (for example, `webshop_*.log`).
- File formats:
 - Fixed length
 - Delimited
 - Customized format
- Loading options:
 - Immediate or deferred loading
 - Overwrite or append
 - Hive external tables

Table 4–4 describes the options for IKM File to Hive (Load Data). See the knowledge module for additional details.

Table 4–4 *IKM File to Hive Options*

Option	Description
CREATE_TARG_TABLE	Create target table.
TRUNCATE	Truncate data in target table.
FILE_IS_LOCAL	Is the file in the local file system or in HDFS?
EXTERNAL_TABLE	Use an externally managed Hive table.
USE_STAGING_TABLE	Use a Hive staging table. Select this option if the source and target do not match or if the partition column value is part of the data file. If the partitioning value is provided by a file name fragment or a constant in target mapping, then set this value to <code>false</code> .
DELETE_TEMPORARY_OBJECTS	Remove temporary objects after the interface execution.
DEFER_TARGET_LOAD	Load data into the final target now or defer?
OVERRIDE_ROW_FORMAT	Provide a parsing expression for handling a custom file format to perform the mapping from source to target.
STOP_ON_FILE_NOT_FOUND	Stop if no source file is found?

Validating and Transforming Data Within Hive

After loading data into Hive, you can validate and transform the data using the following knowledge modules.

IKM Hive Control Append

This knowledge module validates and controls the data, and integrates it into a Hive target table in truncate/insert (append) mode. Invalid data is isolated in an error table and can be recycled. IKM Hive Control Append supports inline view interfaces that use either this knowledge module or IKM Hive Transform.

[Table 4–5](#) lists the options. See the knowledge module for additional details.

Table 4–5 *IKM Hive Control Append Options*

Option	Description
FLOW_CONTROL	Validate incoming data?
RECYCLE_ERRORS	Reintegrate data from error table?
STATIC_CONTROL	Validate data after load?
CREATE_TARG_TABLE	Create target table?
DELETE_TEMPORARY_OBJECTS	Remove temporary objects after execution?
TRUNCATE	Truncate data in target table?

CKM Hive

This knowledge module checks data integrity for Hive tables. It verifies the validity of the constraints of a Hive data store and diverts the invalid records to an error table. You can use CKM Hive for static control and flow control. You must also define these constraints on the stored data.

[Table 4–6](#) lists the options for this check knowledge module. See the knowledge module for additional details.

Table 4–6 *CKM Hive Options*

Option	Description
DROP_ERROR_TABLE	Drop error table before execution?

IKM Hive Transform

This knowledge module performs transformations. It uses a shell script to transform the data, and then integrates it into a Hive target table using replace mode. The knowledge module supports inline view interfaces and can be used as an inline-view for IKM Hive Control Append.

The transformation script must read the input columns in the order defined by the source data store. Only mapped source columns are streamed into the transformations. The transformation script must provide the output columns in the order defined by the target data store.

[Table 4–7](#) lists the options for this integration knowledge module. See the knowledge module for additional details.

Table 4–7 *IKM Hive Transform Options*

Option	Description
CREATE_TARG_TABLE	Create target table?
DELETE_TEMPORARY_OBJECTS	Remove temporary objects after execution?
TRANSFORM_SCRIPT_NAME	Script file name
TRANSFORM_SCRIPT	Script content

Table 4–7 (Cont.) IKM Hive Transform Options

Option	Description
PRE_TRANSFORM_DISTRIBUTE	Provides an optional, comma-separated list of source column names, which enables the knowledge module to distribute the data before the transformation script is applied
PRE_TRANSFORM_SORT	Provide an optional, comma-separated list of source column names, which enables the knowledge module to sort the data before the transformation script is applied
POST_TRANSFORM_DISTRIBUTE	Provides an optional, comma-separated list of target column names, which enables the knowledge module to distribute the data after the transformation script is applied
POST_TRANSFORM_SORT	Provides an optional, comma-separated list of target column names, which enables the knowledge module to sort the data after the transformation script is applied

Loading Data into an Oracle Database from Hive and HDFS

IKM File-Hive to Oracle (OLH) integrates data from an HDFS file or Hive source into an Oracle database target using Oracle Loader for Hadoop. Using the interface configuration and the selected options, the knowledge module generates an appropriate Oracle Database target instance. Hive and Hadoop versions must follow the Oracle Loader for Hadoop requirements.

See Also:

- ["Oracle Loader for Hadoop Setup"](#) on page 1-6 for required versions of Hadoop and Hive
- ["Setting Up the Oracle Data Integrator Agent to Execute Hadoop Jobs"](#) on page 4-4 for required environment variable settings

[Table 4–8](#) lists the options for this integration knowledge module. See the knowledge module for additional details.

Table 4–8 IKM File - Hive to Oracle (OLH)

Option	Description
OLH_OUTPUT_MODE	Specify JDBC, OCI, or Data Pump for data transfer.
CREATE_TARG_TABLE	Create target table?
USE_HIVE_STAGING_TABLE	Materialize Hive source data before extract?
USE_ORACLE_STAGING_TABLE	Use an Oracle database staging table?
EXT_TAB_DIR_LOCATION	Shared file path used for Oracle Data Pump transfer.
TEMP_DIR	Local path for temporary files.
MAPRED_OUTPUT_BASE_DIR	HDFS directory for Oracle Loader for Hadoop output files.
FLOW_TABLE_OPTIONS	Options for flow (stage) table creation when you are using an Oracle database staging table.
DELETE_TEMPORARY_OBJECTS	Remove temporary objects after execution?
OVERRIDE_INPUTFORMAT	Set to handle custom file formats.
EXTRA_OLH_CONF_PROPERTIES	Optional Oracle Loader for Hadoop configuration file properties
TRUNCATE	Truncate data in target table?

Table 4–8 (Cont.) IKM File - Hive to Oracle (OLH)

Option	Description
DELETE_ALL	Delete all data in target table?

Oracle R Connector for Hadoop

This chapter describes R support for big data. It contains the following sections:

- [About Oracle R Connector for Hadoop](#)
- [Scenarios for Using Oracle R Packages](#)
- [Security Notes for Oracle R Connector for Hadoop](#)
- [Functions in Alphabetical Order](#)
- [Functions by Category](#)
- [ORCH mapred.config Class](#)
- [Example R Programs Using Oracle R Connector for Hadoop](#)

About Oracle R Connector for Hadoop

Oracle R Connector for Hadoop is an R package that provides an interface between the local R environment and Apache Hadoop. You install and load this package as you would any other R package. Using simple R functions, you can copy data between R memory, the local file system, and HDFS. You can schedule R programs to execute as Hadoop MapReduce jobs and return the results to any of those locations.

To use Oracle R Connector for Hadoop, you should be familiar with R programming and statistical methods.

See Also: To get started using R, see the R Project website at

<http://www.r-project.org/>

Oracle R Connector for Hadoop APIs

Oracle R Connector for Hadoop provides access from a local R client to Apache Hadoop using these APIs:

- `hadoop`: Provides an interface to Hadoop MapReduce.
- `hdfs`: Provides an interface to HDFS
- `orch`: Provides an interface between the local R instance and Oracle Database

All of the APIs are included in the ORCH library. The functions are listed in this chapter in alphabetical order.

Access to Oracle Database

A separate product, Oracle R Enterprise, provides direct read and write access to Oracle Database and enables you to perform statistical analysis on database tables, views, and other data objects. Access to the data stored in an Oracle database is always restricted to the access rights granted by your DBA.

After analyzing the unstructured data in HDFS using Oracle R Connector for Hadoop, you can store the results in an Oracle database using Oracle R Enterprise. You can then perform additional analysis on this smaller set of data.

Oracle R Enterprise is included in the Oracle Advanced Analytics option to Oracle Database Enterprise Edition. It is not one of the Oracle Big Data Connectors.

See Also: *Oracle R Enterprise User's Guide*

Scenarios for Using Oracle R Packages

The following scenario may help you identify opportunities for using Oracle R Connector for Hadoop with Oracle R Enterprise.

Using Oracle R Connector for Hadoop, you might look for files that you have access to on HDFS and then schedule R calculations to execute on data in one such file. Furthermore, you can upload data stored in text files on your local file system into HDFS for calculations, schedule an R script for execution on the Hadoop cluster, and download the results into a local file.

Using Oracle R Enterprise, you can open the R interface and connect to Oracle Database to work on the tables and views that are visible based on your database privileges. You can filter out rows, add derived columns, project new columns, and perform visual and statistical analysis.

Again using Oracle R Connector for Hadoop, you might deploy a MapReduce job on Hadoop for CPU-intensive calculations written in R. The calculation can use data stored in HDFS or, with Oracle R Enterprise, in an Oracle database. You can return the output of the calculation to an Oracle database and to the R console for visualization or additional processing.

Security Notes for Oracle R Connector for Hadoop

Oracle R Connector for Hadoop invokes the Sqoop utility to connect to Oracle Database either to extract data or to store results. **Sqoop** is a command-line utility for Hadoop that imports and exports data between HDFS or Hive and structured databases. The name Sqoop comes from "SQL to Hadoop."

The following explains how Oracle R Connector for Hadoop stores a database user password and sends it to Sqoop.

Oracle R Connector for Hadoop stores a user password only when the user establishes the database connection in a mode that does not require reentering the password each time. The password is stored encrypted in memory. See [orch.connect](#) on page 5-46.

Oracle R Connector for Hadoop generates a configuration file for Sqoop and uses it to invoke Sqoop locally. The file contains the user's database password obtained by either prompting the user or from the encrypted in-memory representation. The file has local user access permissions only. The file is created, the permissions are set explicitly, and then the file is open for writing and filled with data.

Sqoop uses the configuration file to generate custom JAR files dynamically for the specific database job and passes the JAR files to the Hadoop client software. The password is stored inside the compiled JAR file; it is not stored in plain text.

The JAR file is transferred to the Hadoop cluster over a network connection. The network connection and the transfer protocol are specific to Hadoop, such as port 5900.

The configuration file is deleted after Sqoop finishes compiling its JAR files and starts its own Hadoop jobs.

Functions in Alphabetical Order

```
hadoop.exec  
hadoop.run  
hdfs.attach  
hdfs.cd  
hdfs.cp  
hdfs.describe  
hdfs.download  
hdfs.exists  
hdfs.get  
hdfs.id  
hdfs.ls  
hdfs.mkdir  
hdfs.mv  
hdfs.parts  
hdfs.pull  
hdfs.push  
hdfs.put  
hdfs.pwd  
hdfs.rm  
hdfs.rmdir  
hdfs.root  
hdfs.sample  
hdfs.setroot  
hdfs.size  
hdfs.upload  
is.hdfs.id  
orch.connect  
orch.dbcon  
orch.dbg.off  
orch.dbg.on  
orch.dbg.output  
orch.dbinfo  
orch.disconnect  
orch.dryrun  
orch.export  
orch.keyval  
orch.keyvals  
orch.pack  
orch.reconnect  
orch.unpack  
orch.version
```

Functions by Category

The functions are grouped into these categories:

- [Making Connections](#)
- [Copying Data](#)
- [Exploring Files](#)
- [Writing MapReduce Functions](#)
- [Debugging Scripts](#)
- [Executing Scripts](#)

Making Connections

```
orch.connect  
orch.dbcon  
orch.dbinfo  
orch.disconnect  
orch.reconnect
```

Copying Data

```
hdfs.attach  
hdfs.cp  
hdfs.download  
hdfs.get  
hdfs.mv  
hdfs.pull  
hdfs.push  
hdfs.put  
hdfs.upload  
orch.export  
orch.pack  
orch.unpack
```

Exploring Files

```
hdfs.cd  
hdfs.describe  
hdfs.exists  
hdfs.id  
hdfs.ls  
hdfs.mkdir  
hdfs.parts  
hdfs.pwd  
hdfs.rm  
hdfs.rmdir  
hdfs.root  
hdfs.sample  
hdfs.setroot  
hdfs.size  
is.hdfs.id
```

Writing MapReduce Functions

```
orch.dryrun  
orch.keyval  
orch.keyvals
```

Debugging Scripts

```
orch.dbg.off  
orch.dbg.on  
orch.dbg.output  
orch.version
```

Executing Scripts

```
hadoop.exec  
hadoop.run  
orch.dryrun
```

ORCH mapred.config Class

The `hadoop.exec` and `hadoop.run` functions have an optional argument, `config`, for configuring the resultant MapReduce job. This argument is an instance of the `mapred.config` class.

The `mapred.config` class has these slots:

hdfs.access

Set to `TRUE` to allow access to the HDFS.* functions in the mappers, reducers, and combiners, or set to `FALSE` to restrict access (default).

job.name

A descriptive name for the job so that you can monitor its progress more easily.

map.input

The mapper input data type: `data.frame`, `list`, or `vector` (default).

map.output

The name of a sample data frame that the mapper generates, which defines the output structure from the mappers. This data frame helps the reducers to decipher the metadata of the files generated by the mappers.

map.split

The number of data values given at one time to the mapper:

- `0` sends all the values at one time.
- `1` sends one row only to the mapper at a time.
- `n` sends a minimum of `n` rows to the mapper at a time. In this syntax, `n` is an integer greater than 1. Some algorithms require a minimum number of rows to function.

map.tasks

The number of mappers to run. Specify `1` to run the mappers sequentially; specify a larger integer to run the mappers in parallel.

map.valkey

Set to `TRUE` to duplicate the keys as data values for the mapper, or `FALSE` to use the keys only as keys (default).

min.split.size

The minimum number of rows to send to the mapper at one time.

reduce.input

A reducer input data type: `data.frame` or `list` (default).

reduce.output

The name of a sample data frame that defines the output structure from the reducers. The reducer generates this data frame, which is used to interpret the metadata of the final file it generates.

reduce.split

The number of data values given at one time to the reducer. See the values for [map.split](#).

reduce.tasks

The number of reducers to run. Specify 1 to run the reducers sequentially; specify a larger integer to run the reducers in parallel.

reduce.valkey

Set to `TRUE` to duplicate the keys as data values for the reducer, or `FALSE` to use the keys only as keys (default).

verbose

Set to `TRUE` to generate diagnostic information, or `FALSE` otherwise.

Example R Programs Using Oracle R Connector for Hadoop

The ORCH package includes sample code to help you learn to adapt your R programs to run on a Hadoop cluster using Oracle R Connector for Hadoop. This topic describes these examples and demonstrations.

Using the Examples

The examples show how you use the Oracle R Connector for Hadoop API. You can view them in a text editor after extracting them from `orch.tgz`. They are located in the `orch/inst/examples` directory.

Example R Programs

example-debug.R

Shows how to use key-value pairs, enable debugging, and check for errors. The mapper function selects only the SFO airport data from the `ONTIME_S` data set, and the reducer calculates the mean arrival delay.

This program requires Oracle R Enterprise for the `ore.pull` function. It uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.dbg.on  
orch.dbg.output  
orch.keyval
```

example-filter1.R

Shows how to use key-value pairs. The mapper function selects cars with a distance value greater than 30 from the `cars` data set, and the reducer function calculates the mean distance for each speed.

This program uses the following ORCH functions:


```
hadoop.run  
hdfs.get  
hdfs.put  
orch.keyval
```

example-filter2.R

Shows how to use values only. The mapper function selects cars with a distance greater than 30 and a speed greater than 14 from the `cars` data set, and the reducer function calculates the mean speed and distance as one value pair.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.keyval
```

example-filter3.R

Shows how to load a local file into HDFS. The mapper and reducer functions are the same as [example-filter2.R](#).

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.download  
hdfs.upload  
orch.keyval
```

example-group.apply.R

Shows how to build a parallel model and generate a graph. The mapper partitions the data based on the petal lengths in the `iris` data set, and the reducer uses basic R statistics and graphics functions to fit the data into a linear model and plot a graph.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.download  
hdfs.exists  
hdfs.get  
hdfs.id  
hdfs.mkdir  
hdfs.put  
hdfs.rmdir  
hdfs.upload  
orch.pack  
orch.export  
orch.unpack
```

example-kmeans.R

Shows how to run a simple logistic regression in Hadoop. This program defines a k-means clustering function and generates random points for a clustering test. The results are printed or graphed.

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.get  
hdfs.put  
orch.export
```

example-lm.R

Shows how to define multiple mappers and one reducer that merges all results. The program calculates a linear regression using the `iris` data set.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.export  
orch.pack  
orch.unpack
```

example-lmq.R

Shows more complex analysis, mappers, and reducers for calculating a linear regression using the `iris` data set. The values are computed in three phases by executing three MapReduce jobs.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.export  
orch.pack  
orch.unpack
```

example-logreg.R

Performs a one-dimensional, logistic regression on the `cars` data set.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.put  
orch.export
```

example-map.df.R

Shows how to run the mapper with an unlimited number of records at one time input as a data frame. The mapper selects cars with a distance greater than 30 from the `cars` data set and calculates the mean distance. The reducer merges the results.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put
```

example-map.list.R

Shows how to run the mapper with an unlimited number of records at one time input as a list. The mapper selects cars with a distance greater than 30 from the `cars` data set and calculates the mean distance. The reducer merges the results.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put
```

example-model.plot.R

Shows how to create models and graphs using HDFS. The mapper provides key-value pairs from the `iris` data set to the reducer. The reducer creates a linear model from data extracted from the data set, plots the results, and saves them in an HDFS file.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.download  
hdfs.exists  
hdfs.get  
hdfs.id  
hdfs.mkdir  
hdfs.put  
hdfs.rmdir  
hdfs.upload  
orch.export  
orch.pack
```

example-model.prep.R

Shows how to distribute data across several map tasks. The mapper generates a data frame from a slice of input data from the `iris` data set. The reducer merges the data frames into one output data set.

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.get  
hdfs.put  
orch.export  
orch.keyvals
```

example-rlm.R

Shows how to convert a simple R program into one that can run as a MapReduce job on a Hadoop cluster. In this example, the program calculates and graphs a linear model on the `cars` data set using basic R functions.

This program uses the following ORCH functions:

```
hadoop.run  
orch.keyvals  
orch.unpack
```

example-split.map.R

Shows how to split the data in the mapper. The first job runs the mapper in list mode and splits the list in the mapper. The second job splits a data frame in the mapper. Both jobs use the `cars` data set.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.keyval
```

example-split.reduce.R

Shows how to split the data from the `cars` data set in the reducer.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.get  
hdfs.put  
orch.keyval
```

example-sum.R

Shows how to perform a sum operation in a MapReduce job. The first job sums a vector of numeric values, and the second job sums all columns of a data frame.

This program uses the following ORCH functions:

```
hadoop.run  
orch.keyval
```

example-teragen.matrix.R

Shows how to generate large data sets in a matrix for testing programs in Hadoop. The mappers generate samples of random data, and the reducers merge them.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.put  
orch.export  
orch.keyvals
```

example-teragen.xy.R

Shows how to generate large data sets in a data frame for testing programs in Hadoop. The mappers generate samples of random data, and the reducers merge them.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.put  
orch.export  
orch.keyvals
```

example-teragen2.xy.R

Shows how to generate large data sets in a data frame for testing programs in Hadoop. One mapper generates small samples of random data, and the reducers merge them.

This program uses the following ORCH functions:

```
hadoop.run  
hdfs.put  
orch.export  
orch.keyvals
```

example-terasort.R

Provides an example of a TeraSort job on a set of randomly generated values.

This program uses the following ORCH function:

```
hadoop.run
```

Using the Demos

The demos illustrate various analytic techniques that use the Oracle R Connector for Hadoop API. You can view them in a text editor after extracting them from `orch.tgz`. They are located in the `orch/inst/demos` directory.

Demo R Programs

demo-bagged.clust.R

Provides an example of bagged clustering using randomly generated values. The mappers perform k-means clustering analysis on a subset of the data and generate centroids for the reducers. The reducers combine the centroids into a hierarchical cluster and store the `hclust` object in HDFS.

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.put
```

```
is.hdfs.id  
orch.export  
orch.keyval  
orch.pack  
orch.unpack
```

demo-distance.R

Calculates a similarity matrix using a Euclidian distance metric. This program uses the MovieLens data set, which you can download from

<http://grouplens.org/node/73/>

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.put  
hdfs.sample  
is.hdfs.id  
orch.keyval  
orch.pack  
orch.unpack
```

demo-kmeans.R

Performs k-means clustering. You must provide this demo with two parameters: an `ore-frame` object and a matrix of cluster centers. You must have Oracle R Enterprise to create an `ore-frame` object.

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.get  
orch.export  
orch.keyval  
orch.pack  
orch.unpack
```

demo-pca.R

Calculates the statistics for the principal components analysis (PCA) of a data set. You must provide this demo with a parameter for the file name of the data set. This program creates a tree of multiple MapReduce jobs, which reduce the number of records at each stage. The final MapReduce job performs the final merge operation and generates the statistics for the entire data set.

This program uses the following ORCH functions:

```
hadoop.run  
orch.export  
orch.keyval  
orch.pack  
orch.unpack
```

demo-pearson.R

Calculates a similarity matrix using Pearson's correlation metric. This program uses the MovieLens data set, which you can download from

<http://grouplens.org/node/73/>

This program uses the following ORCH functions:

```
hadoop.exec  
hdfs.put  
is.hdfs.id  
orch.keyval
```

```
orch.pack  
orch.unpack
```

hadoop.exec

Starts the Hadoop engine and sends the mapper, reducer, and combiner R functions for execution. You must load the data into HDFS first.

Usage

```
hadoop.exec (  
  dfs.id,  
  mapper,  
  reducer,  
  combiner,  
  export,  
  init,  
  final,  
  job.name,  
  config)
```

Arguments

dfs.id

The name of a file in HDFS containing data to be processed. The file name can include a path that is either absolute or relative to the current path.

mapper

Name of a mapper function written in the R language.

reducer

Name of a reducer function written in the R language (optional).

combiner

Name of a combiner function written in the R language (optional).

export

Names of exported R objects from your current R environment that are referenced by any of the mapper, reducer, or combiner functions (optional).

init

A function that is executed once before the mapper function begins (optional).

final

A function that is executed once after the reducer function completes (optional).

job.name

A descriptive name that you can use to track the progress of the MapReduce job instead of the automatically generated job name (optional).

config

Sets the configuration parameters for the MapReduce job (optional).

This argument is an instance of the `mapred.config` class, and thus it has this format:

```
config = new("mapred.config", param1, param2, ...)
```

See "[ORCH mapred.config Class](#)" on page 5-5 for a description of this class.

Usage Notes

This function provides more control of the data flow than the [hadoop.run](#) function. You must use `hadoop.exec` when chaining several mappers and reducers in a pipeline, because the data does not leave HDFS. The results are stored in HDFS files.

Return Value

Data object identifier in HDFS

See Also

[hadoop.run](#) on page 5-16, [orch.dryrun](#) on page 5-56

Example

This sample script uses `hdfs.attach` to obtain the object identifier of a small, sample data file in HDFS named `ontime_R`.

The MapReduce function counts the number of on-time flights arriving in the San Francisco International Airport (SFO).

```
dfs <- hdfs.attach('ontime_R')
res <- NULL
res <- hadoop.exec(
  dfs,
  mapper = function(key, ontime) {
    if (key == 'SFO') {
      keyval(key, ontime)
    }
  },
  reducer = function(key, vals) {
    sumAD <- 0
    count <- 0
    for (x in vals) {
      if (!is.na(x$ARRDELAY)) {sumAD <- sumAD + x$ARRDELAY; count <- count +
1}
    }
    res <- sumAD / count
    keyval(key, res)
  }
)
```

After the script runs, the `res` variable identifies the location of the results in an HDFS file named `/user/oracle/xq/orch3d0b8218`:

```
R> res
[1] "/user/oracle/xq/orch3d0b8218"
attr(,"dfs.id")
[1] TRUE
R> print(hdfs.get(res))
  val1    val2
1 SFO 27.05804
```

This code fragment is extracted from [example-kmeans.R](#). The `export` option identifies the location of the `ncenters` generated data set, which is exported as an HDFS file. The `config` options provide a MapReduce job name of `k-means.1`, and the mapper output format of a data frame.

```
mapf <- data.frame(key=0, val1=0, val2=0)
dfs.points <- hdfs.put(points)
dfs.centers <- hadoop.exec(
```

```
dfs.id = dfs.points,
mapper = function(k,v) {
  keyval(sample(1:ncenters,1), v)
},
reducer = function(k,vv) {
  vv <- sapply(vv, unlist)
  keyval(NULL, c(mean(vv[1,]), mean(vv[2,])))
},
export = orch.export(ncenters),
config = new("mapred.config",
  job.name = "k-means.1",
  map.output = mapf)
```

hadoop.run

Starts the Hadoop engine and sends the mapper, reducer, and combiner R functions for execution. If the data is not already stored in HDFS, then `hadoop.run` first copies the data there.

Usage

```
hadoop.run(  
  data,  
  mapper,  
  reducer,  
  combiner,  
  export,  
  init,  
  final,  
  job.name,  
  config)
```

Arguments

data

Data frame, Oracle R Enterprise frame (`ore.frame`), or an HDFS file descriptor.

mapper

Name of a mapper function written in the R language.

reducer

Name of a reducer function written in the R language (optional).

combiner

Name of a combiner function written in the R language (optional).

export

Names of exported R objects.

init

A function that is executed once before the mapper function begins (optional).

final

A function that is executed once after the reducer function completes (optional).

job.name

A descriptive name that you can use to track the progress of the job instead of the automatically generated job name (optional).

config

Sets the configuration parameters for the MapReduce job (optional).

This argument is an instance of the `mapred.config` class, and so it has this format:

```
config = new("mapred.config", param1, param2, ...)
```

See "[ORCH mapred.config Class](#)" on page 5-5 for a description of this class.

Usage Notes

The `hadoop.run` function returns the results from HDFS to the source of the input data. For example, the results for HDFS input data are kept in HDFS, and the results for `ore.frame` input data are copied into an Oracle database.

Return Value

An object in the same format as the input data

See Also

[hadoop.exec](#) on page 5-13, [orch.dryrun](#) on page 5-56

Example

This sample script uses `hdfs.attach` to obtain the object identifier of a small, sample data file in HDFS named `ontime_R`.

The MapReduce function counts the number of on-time flights arriving in the San Francisco International Airport (SFO).

```
dfs <- hdfs.attach('ontime_R')
res <- NULL
res <- hadoop.run(
  dfs,
  mapper = function(key, ontime) {
    if (key == 'SFO') {
      keyval(key, ontime)
    }
  },
  reducer = function(key, vals) {
    sumAD <- 0
    count <- 0
    for (x in vals) {
      if (!is.na(x$ARRDELAY)) {sumAD <- sumAD + x$ARRDELAY; count <- count +
1}
    }
    res <- sumAD / count
    keyval(key, res)
  }
)
```

After the script runs, the location of the results is identified by the `res` variable, in an HDFS file named `/user/oracle/xq/orch3d0b8218`:

```
R> res
[1] "/user/oracle/xq/orch3d0b8218"
attr(,"dfs.id")
[1] TRUE
R> print(hdfs.get(res))
  val1    val2
1 SFO 27.05804
```

hdfs.attach

Copies data from an unstructured data file in HDFS into the Oracle R Connector for Hadoop framework. By default, data files in HDFS are not visible to the connector. However, if you know the name of the data file, you can use this function to attach it to the Oracle R Connector for Hadoop name space.

Usage

```
hdfs.attach(  
  dfs.name,  
  force)
```

Arguments

dfs.name

The name of a file in HDFS.

force

Controls whether the function attempts to discover the structure of the file and the data type of each column.

`FALSE` for comma-separated value (CSV) files (default). If a file does not have metadata identifying the names and data types of the columns, then the function samples the data to deduce the data type as number or string. It then re-creates the file with the appropriate metadata.

`TRUE` for non-CVS files, including binary files. This setting prevents the function from trying to discover the metadata; instead, it simply attaches the file.

Usage Notes

Use this function to attach a CSV file to your R environment, just as you might attach a data frame.

Oracle R Connector for Hadoop does not support the processing of attached non-CVS files. Nonetheless, you can attach a non-CVS file, download it to your local computer, and use it as desired. Alternatively, you can attach the file for use as input to a Hadoop application.

Return Value

The object ID of the file in HDFS, or `NULL` if the operation failed

See Also

[hdfs.download](#) on page 5-23

Example

This example stores the object ID of `ontime_R` in a variable named `dfs`, and then displays its value.

```
R> dfs <- hdfs.attach('ontime_R')  
R> dfs  
[1] "/user/oracle/xq/ontime_R"  
attr(,"dfs.id")
```

[1] TRUE

hdfs.cd

Sets the default HDFS path.

Usage

```
hdfs.cd(dfs.path)
```

Arguments

dfs.path

A path that is either absolute or relative to the current path.

Return Value

TRUE if the path is changed successfully, or FALSE if the operation failed

Example

This example changes the current directory from /user/oracle to /user/oracle/sample:

```
R> hdfs.cd("sample")  
[1] "/user/oracle/sample"
```

hdfs.cp

Copies an HDFS file from one location to another.

Usage

```
hdfs.cp(  
  dfs.src,  
  dfs.dst,  
  force)
```

Arguments

dfs.src

The name of the source file to be copied. The file name can include a path that is either absolute or relative to the current path.

dfs.dst

The name of the copied file. The file name can include a path that is either absolute or relative to the current path.

force

Set to `TRUE` to overwrite an existing file, or set to `FALSE` to display an error message (default).

Return Value

`NULL` for a successful copy, or `FALSE` for a failed attempt

Example

This example copies a file named `weblog` in the parent directory and overwrites the existing `weblog` file:

```
R> hdfs.cp("weblog", "..", force=T)
```

hdfs.describe

Returns the metadata associated with a file in HDFS.

Usage

```
hdfs.describe(dfs.id)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

Return Value

A data frame containing the metadata, or NULL if no metadata was available in HDFS

Example

This example provides information about an HDFS file named `ontime_DB`:

```
R> hdfs.describe('ontime_DB')
      name
1      path
2      origin
3      class
4      types
5      dim
6      names
7      has.key
8      key.column
9      null.key
10     has.rownames
11     size
12     parts

      values
2      ontime_DB
3      unknown
.
.
.
```

hdfs.download

Copies a file from HDFS to the local file system.

Usage

```
hdfs.download(  
  dfs.id,  
  filename,  
  overwrite)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

filename

The name of a file in the local file system where the data is copied.

overwrite

Controls whether the operation can overwrite an existing local file. Set to TRUE to overwrite *filename*, or FALSE to signal an error (default).

Usage Notes

This function provides the fastest and easiest way to copy a file from HDFS. No data transformations occur except merging multiple parts into a single file. The local file has the exact same data as the HDFS file.

Return Value

Local file name, or NULL if the copy failed

Example

This example displays a list of files in the current HDFS directory and copies ontime2000.DB to the local file system as /home/oracle/ontime2000.dat.

```
R> hdfs.ls()  
[1] "ontime2000_DB" "ontime_DB"      "ontime_File"  "ontime_R"     "testdata.dat"  
R> tmpfile <- hdfs.download("ontime2000_DB", "/home/oracle/ontime2000.dat",  
  overwrite=F)  
R> tmpfile  
[1] "/home/oracle/ontime2000.dat"
```

hdfs.exists

Verifies that a file exists in HDFS.

Usage

```
hdfs.exists(  
  dfs.id)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

Usage Notes

If this function returns `TRUE`, then you can attach the data and use it in a [hadoop.run](#) function. You can also use this function to validate an HDFS identifier and ensure that the data exists.

Return Value

`TRUE` if the identifier is valid and the data exists, or `FALSE` if the object is not found

See Also

[is.hdfs.id](#) on page 5-45

Example

This example shows that the `ontime_R` file exists.

```
R> hdfs.exists("ontime_R")  
[1] TRUE
```

hdfs.get

Copies data from HDFS into a data frame in the local R environment. All metadata is extracted and all attributes, such as column names and data types, are restored if the data originated in an R environment. Otherwise, generic attributes like `val1` and `val2` are assigned.

Usage

```
hdfs.get(  
  dfs.id,  
  sep)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

sep

The symbol used to separate fields in the file. A comma (,) is the default separator.

Usage Notes

If the HDFS file is small enough to fit into an in-memory R data frame, then you can copy the file using this function instead of the `hdfs.pull` function. The `hdfs.get` function can be faster, because it does not use Sqoop and thus does not have the overhead incurred by `hdfs.pull`.

Return Value

A `data.frame` object in memory in the local R environment pointing to the exported data set, or `NULL` if the operation failed

Example

This example returns the contents of a data frame named `res`.

```
R> print(hdfs.get(res))  
  val1      val2  
1  AA 1361.4643  
2  AS  515.8000  
3  CO 2507.2857  
4  DL 1601.6154  
5  HP  549.4286  
6  NW 2009.7273  
7  TW 1906.0000  
8  UA 1134.0821  
9  US 2387.5000  
10 WN  541.1538
```

hdfs.id

Converts an HDFS path name to an R `dfs.id` object.

Usage

```
hdfs.id(  
  dfs.x,  
  force)
```

Arguments

dfs.x

A string or text expression that resolves to an HDFS file name.

force

Set to `TRUE` if the file need not exist, or set to `FALSE` to ensure that the file does exist.

Return Value

`TRUE` if the string matches an HDFS file name, or `NULL` if a file by that name is not found

Example

This example creates a `dfs.id` object for `/user/oracle/demo`:

```
R> hdfs.id('/user/oracle/demo')  
[1] "user/oracle/demo"  
attr(,"dfs.id")  
[1] TRUE
```

The next example creates a `dfs.id` object named `id` for a nonexistent directory named `/user/oracle/newdemo`, after first failing:

```
R> id<-hdfs.id('/user/oracle/newdemo')  
DBG: 16:11:38 [ER] "/user/oracle/newdemo" is not found  
R> id<-hdfs.id('/user/oracle/newdemo', force=T)  
R> id  
[1] "user/oracle/newdemo"  
attr(,"dfs.id")  
[1] TRUE
```

hdfs.ls

Lists the names of all HDFS directories containing data in the specified path.

Usage

```
hdfs.ls (dfs.path)
```

Arguments

dfs.path

A path relative to the current default path. The default path is the current working directory.

Return Value

A list of data object names in HDFS, or NULL if the specified path is invalid

See Also

[hdfs.cd](#) on page 5-20

Example

This example lists the subdirectories in the current directory:

```
R> hdfs.ls()  
[1] "ontime_DB"   "ontime_FILE" "ontime_R"
```

The next example lists directories in the parent directory:

```
R> hdfs.ls("../")  
[1] "demo"   "input"  "output" "sample" "xq"
```

This example returns NULL because the specified path is not in HDFS.

```
R> hdfs.ls("/bin")  
NULL
```

hdfs.mkdir

Creates a subdirectory in HDFS relative to the current working directory.

Usage

```
hdfs.mkdir(  
  dfs.name,  
  cd)
```

Arguments

dfs.name

Name of the new directory.

cd

TRUE to change the current working directory to the new subdirectory, or FALSE to keep the current working directory (default).

Return Value

Full path of the new directory as a string, or NULL if the directory was not created

Example

This example creates the `/user/oracle/sample` directory.

```
R> hdfs.mkdir('sample', cd=T)  
[1] "/user/oracle/sample"  
attr(,"dfs.path")  
[1] TRUE
```

hdfs.mv

Moves an HDFS file from one location to another.

Usage

```
hdfs.mv(  
  dfs.src,  
  dfs.dst,  
  force)
```

Arguments

dfs.src

The name of the source file to be moved. The file name can include a path that is either absolute or relative to the current path.

dfs.dst

The name of the moved file. The file name can include a path that is either absolute or relative to the current path.

force

Set to `TRUE` to overwrite an existing destination file, or `FALSE` to cancel the operation and display an error message (default).

Return Value

`NULL` for a successful copy, or `FALSE` for a failed attempt

Example

This example moves a file named `weblog` to the `demo` subdirectory and overwrites the existing `weblog` file:

```
R> hdfs.mv("weblog", "./demo", force=T)
```

hdfs.parts

Returns the number of parts composing a file in HDFS.

Usage

```
hdfs.parts(  
  dfs.id)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

Usage Notes

HDFS splits large files into parts, which provide a basis for the parallelization of MapReduce jobs. The more parts an HDFS file has, the more mappers can run in parallel.

Return Value

The number of parts composing the object, or 0 if the object does not exist in HDFS

Example

This example shows that the `ontime_R` file in HDFS has one part:

```
R> hdfs.parts("ontime_R")  
[1] 1
```

hdfs.pull

Copies data from HDFS into an Oracle database.

This operation requires authentication by Oracle Database. See [orch.connect](#) on page 5-46.

Usage

```
hdfs.pull(  
  dfs.id,  
  sep,  
  db.name,  
  overwrite,  
  driver)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

sep

The symbol used to separate fields in the file (optional). A comma (,) is the default separator.

db.name

The name of a table in an Oracle database (optional).

overwrite

Controls whether `db.name` can overwrite a table with the same name. Set to `TRUE` to overwrite the table, or `FALSE` to signal an error (default).

driver

Identifies the driver used to copy the data. This argument is currently ignored because Sqoop is the only supported driver.

Usage Notes

Because this operation is synchronous, copying a large data set may take a while. The prompt reappears and you regain use of R when copying is complete.

To copy large volumes of data into an Oracle database, consider using Oracle Loader for Hadoop. With the Oracle Advanced Analytics option, you can use Oracle R Enterprise to analyze the data.

Return Value

An `ore.frame` object that points to the database table with data loaded from HDFS, or `NULL` if the operation failed

See Also

Oracle R Enterprise User's Guide for a description of `ore.frame` objects.

hdfs.push

Copies data from an Oracle database to HDFS.

This operation requires authentication by Oracle Database. See [orch.connect](#) on page 5-46.

Note: The Oracle R Enterprise library (ORE) must be attached for you to use this function.

Usage

```
hdfs.push(  
  x,  
  key,  
  dfs.name,  
  overwrite,  
  driver,  
  split.by)
```

Arguments

x

An `ore.frame` object with the data in an Oracle database to be pushed.

key

The index or name of the key column.

dfs.name

Unique name for the object in HDFS.

overwrite

TRUE to allow `dfs.name` to overwrite an object with the same name, or FALSE to signal an error (default).

driver

Identifies the driver used to copy the data. This argument is currently ignored because Sqoop is the only supported driver.

split.by

The column to use for data partitioning (optional).

Usage Notes

Because this operation is synchronous, copying a large data set may take a while. The prompt reappears and you regain use of R when copying is complete.

An `ore.frame` object is an Oracle R Enterprise metadata object that points to a database table. It corresponds to an R `data.frame` object.

Return Value

The full path to the file that contains the data set, or NULL if the operation failed

See Also

Oracle R Enterprise User's Guide

Example

This example creates an `ore.frame` object named `ontime_s2000` that contains the rows from the `ONTIME_S` database table in where the year equals 2000. Then `hdfs.push` uses `ontime_s2000` to create `/user/oracle/xq/ontime2000_DB` in HDFS.

```
R> ontime_s2000 <- ONTIME_S[ONTIME_S$YEAR == 2000,]
R> class(ontime_s2000)
[1] "ore.frame"
attr(,"package")
[1] "OREbase"
R> ontime2000.dfs <- hdfs.push(ontime_s2000, key='DEST', dfs.name='ontime2000_DB')
R> ontime2000.dfs
[1] "/user/oracle/xq/ontime2000_DB"
attr(,"dfs.id")
[1] TRUE
```

hdfs.put

Copies data from an Oracle database to HDFS. Column names, data types, and other attributes are stored as metadata in HDFS.

Note: The Oracle R Enterprise library (ORE) must be attached for you to use this function.

Usage

```
hdfs.put(  
  data,  
  key,  
  dfs.name,  
  overwrite,  
  rownames)
```

Arguments

data

An `ore.frame` object in the local R environment to be copied to HDFS.

key

The index or name of the key column.

dfs.name

A unique name for the new file.

overwrite

Controls whether `dfs.name` can overwrite a file with the same name. Set to `TRUE` to overwrite the file, or `FALSE` to signal an error.

rownames

Set to `TRUE` to add a sequential number to the beginning of each line of the file, or `FALSE` otherwise.

Usage Notes

You can use `hdfs.put` instead of `hdfs.push` to copy data from `ore.frame` objects, such as database tables, to HDFS. The table must be small enough to fit in R memory; otherwise, the function fails. The `hdfs.put` function first reads all table data into R memory and then transfers it to HDFS. For a small table, this function can be faster than `hdfs.push` because it does not use Sqoop and thus does not have the overhead incurred by `hdfs.push`.

Return Value

The object ID of the new file, or `NULL` if the operation failed

Example

This example creates a file named `/user/oracle/xq/testdata.dat` with the contents of the `dat` data frame.

```
R> myfile <- hdfs.put(dat, key='DEST', dfs.name='testdata.dat')
```

```
R> print(myfile)
[1] "/user/oracle/xq/testdata.dat"
attr(,"dfs.id")
[1] TRUE
```

hdfs.pwd

Identifies the current working directory in HDFS.

Usage

```
hdfs.pwd()
```

Return Value

The current working directory, or NULL if your R environment is not connected to HDFS

Example

This example shows that /user/oracle is the current working directory.

```
R> hdfs.pwd()  
[1] "/user/oracle/"
```

hdfs.rm

Removes a file or directory from HDFS.

Usage

```
hdfs.rm(  
  dfs.id,  
  force)
```

Arguments

dfs.id

The name of a file or directory in HDFS. The name can include a path that is either absolute or relative to the current path.

force

Controls whether a directory that contains files is deleted. Set to `TRUE` to delete the directory and all its files, or `FALSE` to cancel the operation (default).

Usage Notes

All object identifiers in Hadoop pointing to this data are invalid after this operation.

Return Value

`TRUE` if the data is deleted, or `FALSE` if the operation failed

Example

This example removes the file named `data1.log` in the current working HDFS directory:

```
R> hdfs.rm("data1.log")  
[1] TRUE
```

hdfs.rmdir

Deletes a directory in HDFS.

Usage

```
hdfs.rmdir(  
  dfs.name,  
  force)
```

Arguments

dfs.name

Name of the directory in HDFS to delete. The directory can be an absolute path or relative to the current working directory.

force

Controls whether a directory that contains files is deleted. Set to `TRUE` to delete the directory and all its files, or `FALSE` to cancel the operation (default).

Usage Notes

This function deletes all data objects stored in the directory, which invalidates all associated object identifiers in HDFS.

Return Value

`TRUE` if the directory is deleted successfully, or `FALSE` if the operation fails

Example

```
R> hdfs.rmdir("mydata")  
[1] TRUE
```


hdfs.root

Returns the HDFS root directory.

Usage

```
hdfs.root()
```

Return Value

A data frame with the full path of the HDFS root directory

Example

This example identifies `/user/oracle` as the root directory of HDFS.

```
R> hdfs.root()
[1] "/user/oracle"
```

hdfs.sample

Copies a random sample of data from a Hadoop file into an R in-memory object. Use this function to copy a small sample of the original HDFS data for developing the R calculation that you ultimately want to execute on the entire HDFS data set on the Hadoop cluster.

Usage

```
hdfs.sample(  
  dfs.id,  
  lines,  
  sep)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

lines

The number of lines to return as a sample. The default value is 1000 lines.

sep

The symbol used to separate fields in the Hadoop file. A comma (,) is the default separator.

Usage Notes

If the data originated in an R environment, then all metadata is extracted and all attributes are restored, including column names and data types. Otherwise, generic attribute names, like `val1` and `val2`, are assigned.

Return Value

A `data.frame` object with the sample data set, or `NULL` if the operation failed

Example

This example displays the first three lines of the `ontime_R` file.

```
R> hdfs.sample("ontime_R", lines=3)  
  YEAR MONTH MONTH2 DAYOFMONTH DAYOFMONTH2 DAYOFWEEK DEPTIME...  
1 2000   12    NA         31         NA         7    1730...  
2 2000   12    NA         31         NA         7    1752...  
3 2000   12    NA         31         NA         7    1803...
```

hdfs.setroot

Sets the HDFS root directory.

Usage

```
hdfs.setroot(dfs.root)
```

Arguments

dfs.root

The full path of the root directory.

Usage Notes

Use [hdfs.root](#) on page 5-39 to see the current root directory.

Return Value

None

Example

This example changes the HDFS root directory from `/user/oracle` to `/user/oracle/demo`.

```
R> hdfs.root()
[1] "/user/oracle"
R> hdfs.setroot("/user/oracle/demo")
R> hdfs.root()
[1] "/user/oracle/demo"
```

hdfs.size

Returns the size of a file in HDFS.

Usage

```
hdfs.size(  
  dfs.id,  
  units)
```

Arguments

dfs.id

The name of a file in HDFS. The file name can include a path that is either absolute or relative to the current path.

units

Specifies a unit of measurement for the return value:

- KB (kilobytes)
- MB (megabytes)
- GB (gigabytes)
- TB (terabytes)
- PB (petabytes)

The unit defaults to bytes if you omit the argument or enter an unknown value.

Usage Notes

Use this interface to determine, for instance, whether you can copy the contents of an entire HDFS file into local R memory or a local file, or if you can only sample the data while creating a prototype of your R calculation.

Return Value

Size of the object, or 0 if the object does not exist in HDFS

Example

This example returns a file size for ontime_R of 999,839 bytes.

```
R> hdfs.size("ontime_R")  
[1] 999839
```

hdfs.upload

Copies a file from the local file system into HDFS.

Usage

```
hdfs.upload(  
  filename,  
  dfs.name,  
  overwrite,  
  split.size,  
  header)
```

Arguments

filename

Name of a file in the local file system.

dfs.name

Name of the new directory in HDFS.

overwrite

Controls whether `dfs.name` can overwrite a directory with the same name. Set to `TRUE` to overwrite the directory, or `FALSE` to signal an error (default).

split.size

Maximum number of bytes in each part of the Hadoop file (optional).

header

Indicates whether the first line of the local file is a header containing column names. Set to `TRUE` if it has a header, or `FALSE` if it does not (default).

A header enables you to exact the column names and reference the data fields by name instead of by index in your MapReduce R scripts.

Usage Notes

This function provides the fastest and easiest way to copy a file into HDFS. If the file is larger than `split.size`, then Hadoop splits it into two or more parts. The new Hadoop file gets a unique object ID, and each part is named `part-0000x`. Hadoop automatically creates metadata for the file.

Return Value

HDFS object ID for the loaded data, or `NULL` if the copy failed

See Also

[hdfs.download](#) on page 5-23, [hdfs.get](#) on page 5-25, [hdfs.put](#) on page 5-34

Example

This example uploads a file named `ontime_s2000.dat` into HDFS and shows the location of the file, which is stored in a variable named `ontime.dfs_File`.

```
R> ontime.dfs_File <- hdfs.upload('ontime_s2000.dat', dfs.name='ontime_File')  
R> print(ontime.dfs_File)
```

```
[1] "/user/oracle/xq/ontime_File"
```

is.hdfs.id

Indicates whether an R object contains a valid HDFS file identifier.

Usage

```
is.hdfs.id(x)
```

Arguments

x
The name of an R object.

Return Value

TRUE if x is a valid HDFS identifier, or FALSE if it is not

See Also

[hdfs.attach](#) on page 5-18, [hdfs.id](#) on page 5-26

Example

This example shows that `dfs` contains a valid HDFS identifier, which was returned by `hdfs.attach`:

```
R> dfs <- hdfs.attach('ontime_R')
R> is.hdfs.id(dfs)
[1] TRUE
R> print(dfs)
[1] "/user/oracle/xq/ontime_R"
attr(,"dfs.id")
[1] TRUE
```

The next example shows that a valid file name passed as a string is not recognized as a valid file identifier:

```
R> is.hdfs.id('/user/oracle/xq/ontime_R')
[1] FALSE
```

orch.connect

Establishes a connection to Oracle Database.

Usage

```
orch.connect(  
    host,  
    user,  
    sid,  
    passwd,  
    port,  
    secure,  
    driver,  
    silent)
```

Arguments

host

Host name or IP address of the server where Oracle Database is running.

user

Database user name.

sid

System ID (SID) for the Oracle Database instance.

passwd

Password for the database user. If you omit the password, you are prompted for it.

port

Port number for the Oracle Database listener. The default value is 1521.

secure

Authentication setting for Oracle Database:

- **TRUE:** You must enter a database password each time you attempt to connect (default).
- **FALSE:** You must enter a database password only once during a session. The encrypted password is kept in memory and used for subsequent connection attempts.

driver

Driver used to connect to Oracle Database (optional). Sqoop is the default driver.

silent

TRUE to suppress the prompts for missing host, user, password, port, and SID values, or **FALSE** to see them (default).

Usage Notes

Use this function when your analysis requires access to data stored in an Oracle database or to return the results to the database.

With an Oracle Advanced Analytics license for Oracle R Enterprise and a connection to Oracle Database, you can work directly with the data stored in database tables and pass processed data frames to R calculations on Hadoop.

You can reconnect to Oracle Database using the connection object returned by the `orch.dbcon` function.

Return Value

TRUE for a successful and validated connection, or FALSE for a failed connection attempt

See Also

[orch.dbcon](#) on page 5-48, [orch.disconnect](#) on page 5-54

Example

This example installs the ORCH library and connects to Oracle Database on the local system:

```
R> library(ORCH)
Oracle R Connector for Hadoop 0.1.8 (rev.102)
Hadoop 0.20.2-cdh3u3 is up
Sqoop 1.3.0-cdh3u3 is up
R> orch.connect("localhost", "RUSER", "orcl")
Connecting ORCH to RDBMS via [sqoop]
  Host: localhost
  Port: 1521
  SID: orcl
  User: RUSER
Enter password for [RUSER]: password
Connected.
[1] TRUE
```

The next example uses a connection object to reconnect to Oracle Database:

```
R> conn<-orch.dbcon()
R> orch.disconnect()
Disconnected from a database.

R> orch.connect(conn)
Connecting ORCH to RDBMS via [sqoop]
  Host: localhost
  Port: 1521
  SID: orcl
  User: RUSER
Enter password for [RUSER]: password
Connected
[1] TRUE
```

orch.dbcon

Returns a connection object for the current connection to Oracle Database, excluding the authentication credentials.

Usage

```
orch.dbcon()
```

Return Value

A data frame with the connection settings for Oracle Database

Usage Notes

Use the connection object returned by `orch.dbcon` to reconnect to Oracle Database using `orch.connect`.

See Also

[orch.connect](#) on page 5-46

Example

This example shows how you can reconnect to Oracle Database using the connection object returned by `orch.dbcon`:

```
R> orch.connect('localhost', 'RQUSER', 'orcl')
Connecting ORCH to RDBMS via [sqoop]
  Host: localhost
  Port: 1521
  SID: orcl
  User: RQUSER
Enter password for [RQUSER]: password
Connected
[1] TRUE
```

```
R> conn<-orch.dbcon()
R> orch.disconnect()
Disconnected from a database.
```

```
R> orch.connect(conn)
Connecting ORCH to RDBMS via [sqoop]
  Host: localhost
  Port: 1521
  SID: orcl
  User: RQUSER
Enter password for [RQUSER]: password
Connected
[1] TRUE
```

The following shows the connection object returned by `orch.dbcon` in the previous example:

```
R> conn
Object of class "orch.dbcon"
data frame with 0 columns and 0 rows
Slot "ok":
```

```
[1] TRUE

Slot "host":
[1] "localhost"

Slot "port":
[1] 1521

Slot "sid"
[1] "orcl"

Slot "user":
[1] "RQUSER"

Slot "passwd":
[1] ""

Slot "secure":
[1] TRUE

Slot "drv":
[1] "sqoop"
```

orch.dbg.off

Turns off debugging mode.

Usage

```
orch.dbg.off()
```

Return Value

FALSE

See Also

[orch.dbg.on](#) on page 5-51

Example

This example turns off debugging:

```
R> orch.dbg.off()
```

orch.dbg.on

Turns on debugging mode.

Usage

```
orch.dbg.on(severity)
```

Arguments

severity

Identifies the type of error messages that are displayed. You can identify the severity by the number or by the case-insensitive keyword shown in [Table 5-1](#).

Table 5-1 *Debugging Severity Levels*

Keyword	Number	Description
all	11	Return all messages.
critical	1	Return only critical errors.
error	2	Return all errors.
warning	3	Return all warnings.

Return Value

The severity level

See Also

[orch.dbg.output](#) on page 5-52, [orch.dbg.off](#) on page 5-50

Example

This example turns on debugging for all errors:

```
R> severe<-orch.dbg.on(severity<-2)
R> severe
[1] "ERROR" "2"
```

orch.dbg.output

Directs the output from the debugger.

Usage

```
orch.dbg.output (con)
```

Arguments

con

Identifies the stream where the debugging information is sent: `stderr()`, `stdout()`, or a file name.

Usage Notes

You must first turn on debugging mode before redirecting the output.

Return Value

The current stream

See Also

[orch.dbg.on](#) on page 5-51

Example

This example turns on debugging mode and sends the debugging information to `stderr`. The `orch.dbg.output` function returns a description of `stderr`.

```
R> orch.dbg.on('all')
R> err<-orch.dbg.output(stderr())
17:32:11 [SY] debug output set to "stderr"
R> print(err)
description      class      mode      text      opened      can read      can write
  "stderr"       "terminal"  "w"      "text"    "opened"    "no"         "yes"
```

The next example redirects the output to a file named `debug.log`:

```
R> err<-orch.dbg.output('debug.log')
17:37:45 [SY] debug output set to "debug.log"
R> print(err)
[1] "debug.log"
```

orch.dbinfo

Provides information about the current connection.

Usage

```
orch.dbinfo (dbcon)
```

Arguments

dbcon

An Oracle Database connection object.

See Also

[orch.dbcon](#) on page 5-48

orch.disconnect

Disconnects the local R session from Oracle Database.

Usage

```
orch.disconnect(  
  silent,  
  dbcon)
```

Arguments

silent

Set to `TRUE` to suppress all messages, or `FALSE` to see them (default).

dbcon

Set to `TRUE` to display the connection details, or `FALSE` to suppress this information (default).

Usage Notes

No `orch` functions work without a connection to Oracle Database.

Return Value

An Oracle Database connection object when `dbcon` is `TRUE`, otherwise `NULL`

See Also

[orch.connect](#) on page 5-46, [orch.reconnect](#) on page 5-62

Example

This example disconnects the local R session from Oracle Database:

```
R> orch.disconnect()  
Disconnected from a database.
```

The next example disconnects the local R session from Oracle Database and displays the returned connection object:

```
R> oid<-orch.disconnect(silent=TRUE, dbcon=TRUE)  
R> oid  
Object of class "orch.dbcon"  
data frame with 0 columns and 0 rows  
slow "ok":  
[1] TRUE  
  
Slot "host":  
[1] "localhost"  
  
Slot "port":  
[1] 1521  
  
Slot "sid":  
[1] orcl  
  
Slot "user":
```



```
[1] RUSER
```

```
Slot "passwd":
```

```
[1] ""
```

```
Slot "secure":
```

```
[1] TRUE
```

```
Slot "drv":
```

```
[1] "sqoop"
```

orch.dryrun

Switches the execution platform between the local host and the Hadoop cluster. No changes in the R code are required for a dry run.

Usage

```
orch.dryrun (onoff)
```

Arguments

onoff

Set to `TRUE` to run a MapReduce program locally, or `FALSE` to run the program on the Hadoop cluster.

Usage Notes

The `orch.dryrun` function enables you to run a MapReduce program locally on a laptop using a small data set before running it on a Hadoop cluster using a very large data set. The mappers and reducers are run sequentially on row streams from HDFS. The Hadoop cluster is not required for a dry run.

Return Value

The current setting of `orch.dryrun`

See Also

[hadoop.exec](#) on page 5-13, [hadoop.run](#) on page 5-16

Example

This example changes the value of `orch.dryrun` from `FALSE` to `TRUE`.

```
R> orch.dryrun()
[1] FALSE
R> orch.dryrun(onoff<-T)
R> orch.dryrun()
[1] TRUE
```

orch.export

Makes R objects from a user's local R session available in the Hadoop execution environment, so that they can be referenced in MapReduce jobs.

Usage

```
orch.export(...)
```

Arguments

...

One or more variables, data frames, or other in-memory objects, by name or as an explicit definition, in a comma-separated list.

Usage Notes

You can use this function to prepare local variables for use in [hadoop.exec](#) and [hadoop.run](#) functions. The `mapper`, `reducer`, `combiner`, `init`, and `final` arguments can reference the exported variables.

Return Value

A list object

See Also

[hadoop.exec](#) on page 5-13, [hadoop.run](#) on page 5-16

Example

This code fragment shows `orch.export` used in the `export` argument of the `hadoop.run` function:

```
hadoop.run(x,  
  export = orch.export(a=1, b=2),  
  mapper = function(k,v) {  
    x <- a + b  
    orch.keyval(key=NULL, val=x)  
  }  
)
```

The following is a similar code fragment, except that the variables are defined outside the `hadoop.run` function:

```
a=1  
b=2  
hadoop.run(x,  
  export = orch.export(a, b),  
  .  
  .  
  .  
)
```

orch.keyval

Outputs key-value pairs in a MapReduce job.

Usage

```
orch.keyval(  
    key,  
    value)
```

Arguments

key

A scalar value.

value

A data structure such as a scalar, list, data frame, or vector.

Usage Notes

This function can only be used in the mapper, reducer, or combiner arguments of [hadoop.exec](#) and [hadoop.run](#). Because the `orch.keyval` function is not exposed in the ORCH client API, you cannot call it anywhere else.

See Also

[orch.pack](#) on page 5-61

Return Value

(key, value) structures

Example

This code fragment creates a mapper function using `orch.keyval`:

```
hadoop.run(data,  
    mapper = function(k,v) {  
        orch.keyval(k,v)  
    })
```

orch.keyvals

Outputs a set of key-value pairs in a MapReduce job.

Usage

```
orch.keyvals(
  key,
  value)
```

Arguments

key

A scalar value.

value

A data structure such as a scalar, list, data frame, or vector.

Usage Notes

This function can only be used in the mapper, reducer, or combiner arguments of [hadoop.exec](#) and [hadoop.run](#). Because the `orch.keyvals` function is not exposed in the ORCH client API, you cannot call it anywhere else.

See Also

[orch.keyval](#) on page 5-58, [orch.pack](#) on page 5-61

Return Value

(key, value) structures

Example

This code fragment creates a mapper function using `orch.keyval` and a reducer function using `orch.keyvals`:

```
hadoop.run(data,
  mapper(k,v) {
    if (v$value > 10) {
      orch.keyval(k, v)
    }
    else {
      NULL
    }
  },
  reducer(k,vals) {
    orch.keyvals(k,vals)
  }
)
```

The following code fragment shows `orch.keyval` in a for loop to perform the same reduce operation as `orch.keyvals` in the previous example:

```
reducer(k,vals) {
  out <- list()
  for (v in vals) {
```

```
        out <- list(out, orch.keyval(k,vals))
    }
    out
}
```

orch.pack

Compresses one or more in-memory R objects that the mappers or reducers must write as the values in key-value pairs.

Usage

```
orch.pack(...)
```

Arguments

...

One or more variables, data frames, or other in-memory objects in a comma-separated list.

Usage Notes

This function requires the `bitops` package, which you can download from the Comprehensive R Archive Network (CRAN) at

<http://cran.r-project.org/>

You should use this function when passing nonscalar or complex R objects, such as data frames and R classes, between the mapper and reducer functions. You do not need to use it on scalar or other simple objects. You can use `orch.pack` to vary the data formats, data sets, and variable names for each output value.

You should also use `orch.pack` when storing the resultant data set in HDFS. The compressed data set is not corrupted by being stored in an HDFS file.

The `orch.pack` function must always be followed by the `orch.unpack` function to restore the data to a usable format.

Return Value

Compressed character-type data as a long string with no special characters

See Also

[hadoop.exec](#) on page 5-13, [hadoop.run](#) on page 5-16, [orch.keyval](#) on page 5-58, [orch.unpack](#) on page 5-63

Example

This code fragment compresses the content of several R objects into a serialized stream using `orch.pack`, and then creates key-value pairs using `orch.keyval`:

```
orch.keyval(NULL, orch.pack(  
  r = r,  
  qY = qY,  
  YY = YY,  
  nRows = nRows))
```

orch.reconnect

Reconnects to Oracle Database with the credentials previously returned by [orch.disconnect](#).

Note: The `orch.reconnect` function is deprecated in release 1.1 and will be desupported in release 2.0. Use [orch.connect](#) to reconnect using a connection object returned by [orch.dbcon](#).

Usage

```
orch.reconnect (dbcon)
```

Arguments

dbcon

Credentials previously returned by [orch.disconnect](#).

Usage Notes

Oracle R Connector for Hadoop preserves all user credentials and connection attributes, enabling you to reconnect to a previously disconnected session. Depending on the `orch.connect` secure setting for the original connection, you may be prompted for a password. After reconnecting, you can continue data transfer operations between Oracle Database and HDFS.

Reconnecting to a session is faster than opening a new one, because reconnecting does not require extensive connectivity checks.

Return Value

TRUE for a successfully reestablished and validated connection, or FALSE for a failed attempt

See Also

[orch.connect](#) on page 5-46

Example

```
R> orch.reconnect (oid)
Connecting ORCH to RDBMS via [sqoop]
  Host: localhost
  Port: 1521
  SID:  orcl
  User: RQUSER
Enter password for [RQUSER]: password
Connected
[1] TRUE
```

orch.unpack

Restores the R objects that were compressed with a previous call to `orch.pack`.

Usage

```
orch.unpack(...)
```

Arguments

...

The name of a compressed object.

Usage Notes

This function requires the `bitops` package, which you can download from a Comprehensive R Archive Network (CRAN) mirror site at

<http://cran.r-project.org/>

This function is typically used at the beginning of a mapper or reducer function to obtain and prepare the input. However, it can also be used externally, such as at the R console, to unpack the results of a MapReduce job.

Consider this data flow:

ORCH client to mapper to combiner to reducer to ORCH client.

If the data is packed at one stage, then it must be unpacked at the next stage.

Return Value

Uncompressed list-type data, which can contain any number and any type of variables

See Also

[orch.pack](#) on page 5-61

Example

This code fragment restores the data that was compressed in the [orch.pack](#) example:

```
reducer = function(key, vals) {  
  x <-orch.unpack(vals[[1]])  
  r <-x$qy  
  yy <- x$yy  
  nRow <- x$nRows  
  .  
  .  
  .  
}
```

orch.version

Identifies the version of the ORCH package.

Usage

```
orch.version()
```

Return Value

ORCH package version number

Example

This example shows that the ORCH version number is 0.1.8.

```
R> orch.version()  
[1] "0.1.8"
```

A

access drivers
 See drivers
access privileges, Oracle Database, 1-5
ALTER SESSION commands, 2-9
Apache Hadoop distribution, 1-2, 1-3, 1-4, 1-6, 1-8, 3-9
Apache licenses, 3-34, 3-37, 3-38
APPEND hint, 2-9
Avro license, 3-37

B

bagged.clust.R demo, 5-10
balancing loads in Oracle Loader for Hadoop, 3-11
bitops R package, 1-8, 5-61, 5-63

C

cd command
 executing in HDFS from R, 5-20
CDH3 distribution, 1-6
CKM Hive, 4-2, 4-9
client libraries, 1-6
clients
 configuring Hadoop, 1-4, 1-10
combiners
 generating key-value pairs for in R, 5-58, 5-59
 running from R, 5-13, 5-16
Comprehensive R Archive Network
 See CRAN
compressed files, 2-8
compression in R, 5-61
configuration settings
 Hadoop client, 1-4, 1-10
 Oracle Data Integrator agent, 4-4
 Oracle Data Integrator interfaces, 4-7
 Oracle Data Integrator Studio, 4-6
 Oracle Direct Connector for HDFS, 2-5, 2-6
 Oracle Loader for Hadoop, 3-13, 3-15
 Oracle Net, 2-6
 R MapReduce jobs, 5-5
 Sqoop utility, 1-8
configuring a Hadoop client, 1-4, 1-10
connecting to Oracle Database from R, 5-4

connection objects
 reconnecting to Oracle Database with, 5-47
copying data
 from an Oracle database to HDFS using R, 5-32, 5-34
 from HDFS to an Oracle database using R, 5-31
 from HDFS to an R data frame, 5-25
 from local files to HDFS using R, 5-43
copying files
 from HDFS to Linux using R, 5-23
cp command
 executing in HDFS from R, 5-21
CRAN, 5-61, 5-63
CREATE SESSION privilege, 1-5
CREATE TABLE command, 2-2
CREATE_TARG_TABLE option, 4-8, 4-9, 4-10
creating external tables, 2-2
CSV files, 2-1, 2-3, 2-8, 3-2, 3-9, 3-20, 5-18

D

data compression in R, 5-61
data decompression in R, 5-63
data files
 accessing from R, 5-18
 copying database objects into using R, 5-32
 copying into Oracle Database using R, 5-31
 deleting from HDFS using R, 5-37
 obtaining size from R, 5-42
 sampling from R, 5-40
 uploading local, 5-43
 See also files; HDFS files
data frames
 copying HDFS data into, 5-25
 example of mapper output format, 5-14
 example program using, 5-8, 5-9, 5-10
 from Oracle R Enterprise, 5-16
 generated by R mapper functions, 5-5
 generated by R reducer functions, 5-6
 input to mappers, 5-16
data integrity checking, 4-9
Data Pump files, 2-1
 access parameters, 2-2
 generating, 2-1, 3-10
 publishing in Oracle Direct Connector for HDFS, 2-5, 2-8

- specifying location in Oracle Direct Connector for HDFS, 2-3
- Data Pump format, 3-2
- data sources
 - defining for Oracle Data Integrator, 4-3
- data transformations, 4-9
- data types
 - Oracle Loader for Hadoop, 3-21
- data validation in Oracle Data Integrator, 4-2, 4-9
- database
 - copying data from using R, 5-34
- database client installation, 4-5
- database connections
 - disconnecting using R, 5-54
 - from R, 5-46, 5-48
 - querying from R, 5-53
 - reconnecting from R, 5-62
- database directories
 - for external tables, 2-2
 - for Oracle Direct Connector for HDFS, 1-5, 2-2
- database objects
 - copying to HDFS using R, 5-32
- database patches, 1-4, 1-6, 2-1
- database privileges, 1-5
- database system, configuring to run MapReduce jobs, 1-4
- DataServer objects
 - creating for Oracle Data Integrator, 4-3
- debugging in R
 - identifying version, 5-64
 - redirecting output, 5-52
 - turning off, 5-50
 - turning on, 5-51
- debug.R example, 5-6
- decompressing files, 2-8
- decompression in R, 5-63
- DEFER_TARGET_LOAD option, 4-8
- DELETE_ALL option, 4-11
- DELETE_TEMPORARY_OBJECTS option, 4-8, 4-9, 4-10
- delimited text output format
 - Oracle Loader for Hadoop, 3-9
- DelimitedTextInputFormat class, 3-3, 3-18
 - Oracle Loader for Hadoop, 3-3, 3-4
- DelimitedTextOutputFormat class, 3-17, 3-20
- demo code for ORCH, 5-10
- describe command
 - executing in HDFS using R, 5-22
- dfs.id object
 - obtaining in R, 5-26
- directories
 - accessible by Oracle Data Integrator, 4-3
 - creating in HDFS using R, 5-28
 - current HDFS, 5-36
 - for Oracle Loader for Hadoop output, 4-10
 - Hadoop home, 1-4
 - listing in HDFS, 5-27
 - ODI Application Adapter for Hadoop home, 1-7
 - Oracle Data Integrator agent drivers, 4-4
 - Oracle Direct Connector for HDFS home, 1-5

- Oracle Net, 2-6
- R connector examples, 5-6, 5-10
- removing using R, 5-38
- Sqoop home, 1-8
- See also* database directories; root directory
- distance.R demo, 5-11
- downloading software, 1-3, 1-4, 1-7, 1-8, 1-9, 1-10, 5-61, 5-63
- drivers
 - JDBC, 1-8, 3-8, 4-4
 - Oracle Data Integrator agent, 4-4
 - ORACLE_DATAPUMP, 3-2, 3-10, 3-11
 - ORACLE_LOADER, 2-1, 3-2
- DROP_ERROR_TABLE option, 4-9
- dry run execution in R, 5-56

E

- environment variables
 - See* individual variable entries
- example code for ORCH
 - descriptions, 5-6
- exists function in R, 5-24
- exporting R objects to Hadoop, 5-57
- EXT_TAB_DIR_LOCATION option, 4-10
- external tables
 - about, 2-1
 - creating for Oracle Direct Connector for HDFS, 2-2
 - example, 2-3
 - syntax, 2-2
- EXTERNAL_VARIABLE_DATA parameter, 2-2
- EXTERNAL_TABLE option, 4-8
- ExternalTable command, 2-4
 - configuration settings, 2-5
 - example, 2-5
 - path settings, 2-4
 - syntax, 2-4
 - See also* external tables
- EXTRA_OLH_CONF_PROPERTIES option, 4-10

F

- file formats for Oracle Data Integrator, 4-8
- file identifiers
 - validating in R, 5-45
- file size
 - obtaining using R, 5-42
- file sources
 - defining for Oracle Data Integrator, 4-3
- File to Hive KM, 4-2, 4-8
- FILE_IS_LOCAL option, 4-8
- File-Hive to Oracle (OLH) KM, 4-2, 4-10
- files
 - decompressing, 2-8
 - downloading from HDFS, 5-23
 - obtaining HDFS metadata using R, 5-22
 - testing existence in HDFS from R, 5-24
 - uploading local to HDFS, 5-43
 - See also* data files; HDFS files

filter1.R example, 5-6
filter2.R example, 5-7
filter3.R example, 5-7
flex fields, 4-4, 4-7
FLOW_CONTROL option, 4-9
FLOW_TABLE_OPTIONS option, 4-10

G

group.apply.R example, 5-7

H

Hadoop client

configuring, 1-4, 1-10
installing, 1-4

hadoop fs command, 1-4

HADOOP_CLASSPATH

for Oracle Data Integrator, 4-5
for Oracle Loader for Hadoop, 1-6

HADOOP_CLASSPATH environment variable
for Oracle Direct Connector for HDFS, 2-4, 2-5

HADOOP_HOME

for Oracle Data Integrator, 4-4
for Oracle Direct Connector for HDFS, 1-4

HADOOP_HOME environment variable
for Oracle R Connector, 1-9

hadoop.exec function, 5-5

config argument, 5-5
example, 5-7, 5-9, 5-10, 5-11, 5-14
syntax description, 5-13

hadoop.run function, 5-5

config argument, 5-5
example, 5-6, 5-7, 5-8, 5-9, 5-10, 5-11, 5-17
syntax description, 5-16

HDFS client

See Hadoop client

HDFS commands

issuing from R, 5-4

HDFS data

copying in R, 5-4

HDFS directories

acquiring root using R, 5-39
changing in R, 5-20
changing root using R, 5-41
creating a dfs.id object, 5-26
creating in R, 5-28
deleting using R, 5-37, 5-38
identifying current in R, 5-36
listing in R, 5-27
publishing to external table, 2-5

HDFS file identifiers

obtaining in R, 5-26
validating in R, 5-45

HDFS file names

converting to dfs.id object in R, 5-26

HDFS files

access from Oracle Database, 2-1
accessing from R, 5-18
copying data from Oracle Database, 5-32

copying data from the database using R, 5-34
copying data to Oracle Database, 5-31
copying data to R, 5-25
copying from database objects using R, 5-32
copying into Hive, 4-7
copying into Oracle Database using R, 5-31
copying using R, 5-21
creating external tables for accessing, 2-2
deleting using R, 5-37
downloading to local system, 5-23
loading data into an Oracle database, 3-2
moving from R, 5-29
number of parts, 5-30
obtaining size from R, 5-42
sampling from R, 5-40
streaming data into an Oracle database, 2-4
testing existence in HDFS from R, 5-24
uploading local, 5-43

HDFS metadata

obtaining using R, 5-22

HDFS path

changing in R, 5-20

HDFS root directory

changing from R, 5-41
obtaining using R, 5-39

HDFS_BIN_PATH directory, 1-5

hdfs_stream Bash shell script, 1-4

hdfs.attach function

example, 5-14, 5-17, 5-45
syntax description, 5-18

hdfs.cd function

syntax description, 5-20

hdfs.cp function

syntax description, 5-21

hdfs.describe function

syntax example, 5-22

hdfs.download function

example, 5-7, 5-9
syntax description, 5-23

hdfs.exists function

example, 5-7, 5-9
syntax description, 5-24

hdfs.get function

example, 5-6, 5-7, 5-8, 5-9, 5-11
syntax description, 5-25

hdfs.id function

example, 5-7, 5-9
syntax description, 5-26

hdfs.ls function

syntax description, 5-27

hdfs.mkdir function

example, 5-7, 5-9
syntax description, 5-28

hdfs.mv function

syntax description, 5-29

hdfs.parts function

syntax description, 5-30

hdfs.pull function

syntax description, 5-31

hdfs.push function

- syntax description, 5-32
- hdfs.put function
 - example, 5-6, 5-7, 5-8, 5-9, 5-10, 5-11
 - syntax description, 5-34
- hdfs.pwd function
 - syntax description, 5-36
- hdfs.rm function
 - syntax example, 5-37
- hdfs.rmdir function
 - example, 5-7, 5-9
 - syntax description, 5-38
- hdfs.root function
 - syntax description, 5-39
- hdfs.sample function
 - example, 5-11
 - syntax description, 5-40
- hdfs.setroot function
 - syntax description, 5-41
- hdfs.size function
 - syntax description, 5-42
- hdfs.upload function
 - example, 5-7, 5-9
 - syntax description, 5-43
- hints for optimizing queries, 2-9
- Hive application adapters, 4-2, 4-9
- Hive Control Append KM, 4-2, 4-9
- Hive data source for Oracle Data Integrator, 4-3
- Hive database for Oracle Loader for Hadoop, 1-6
- Hive distribution, 1-6
- Hive flex fields, 4-7
- Hive Query Language (HiveQL), 4-2
- Hive tables
 - copying data to Oracle Database, 3-3
 - loading data into (Oracle Data Integrator), 4-7
 - reverse engineering, 4-2, 4-6
 - reverse engineering in Oracle Data Integrator, 4-6
- Hive Transform KM, 4-2, 4-9
- HIVE_HOME, for Oracle Data Integrator, 4-4
- HiveToAvroInputFormat class, 1-6, 3-3

I

- IKM File to Hive, 4-2, 4-8
- IKM File-Hive to Oracle (OLH), 4-2, 4-10
- IKM Hive Control Append, 4-2, 4-9
- IKM Hive Transform, 4-2, 4-9
- IndexedRecord, 3-5
- InputFormat class
 - Oracle Loader for Hadoop, 3-3, 3-4
- INSERT_UPDATE mode, 4-3
- installation
 - Apache Hadoop, 1-4
 - CDH, 1-4
 - Hadoop client, 1-4
 - Oracle Data Integrator Application Adapter for Hadoop, 1-6
 - Oracle Direct Connector for HDFS, 1-3
 - Oracle Loader for Hadoop, 1-6
 - Oracle R Connector for Hadoop, 1-7
 - Sqoop utility, 1-8

- Instant Client libraries, 1-6
- interface configurations
 - Oracle Data Integrator, 4-7
- is.hdfs.id function
 - example, 5-11
 - syntax description, 5-45

J

- JAVA_HOME environment variable
 - for R connector, 1-9
- JAVA_LIBRARY_PATH environment variable
 - for Oracle Data Integrator, 4-5
 - for Oracle Loader for Hadoop, 3-9
- JDBC drivers, 1-8, 3-8, 4-4
 - See* drivers
- job names
 - specifying in Oracle Loader for Hadoop, 3-15
 - specifying in R, 5-5, 5-13, 5-14, 5-16

K

- key-value pairs
 - generating from R objects, 5-58, 5-59
- kmeans.R demo, 5-11
- kmeans.R example, 5-7
- knowledge modules
 - description, 4-2
 - See also* CKM, IKM, and RKM entries

L

- LD_LIBRARY_PATH environment variable
 - for Oracle Data Integrator, 4-5
 - for Oracle Loader for Hadoop, 3-9
- licenses, third-party, 3-34
- lmqr.R example, 5-8
- lm.R example, 5-8
- load balancing
 - in Oracle Loader for Hadoop, 3-11, 3-12
- loadCI, 3-12
- loader map XML schema definition, 3-22
- loaderMap document, 3-5
- loading data files into Hive, 4-7
- loading options for Oracle Data Integrator, 4-8
- local files
 - copying into Hive, 4-7
 - downloading from HDFS, 5-23
- LOG_FILE_NAME option, 4-7
- logreg.R example, 5-8

M

- map.df.R example, 5-8
- map.list.R example, 5-8
- mappers
 - configuration settings, 5-5
 - examples, 5-6
 - generating key-value pairs for in R, 5-58, 5-59
 - running from R, 5-13, 5-16
- MAPRED_OUTPUT_BASE_DIR option, 4-10

- mapred.config class
 - syntax description, 5-5
 - using, 5-13, 5-16
- MapReduce functions
 - writing in R
 - Oracle R Connector for Hadoop
 - MapReduce functions, 5-4
- MapReduce jobs
 - configuring in R, 5-5
 - exporting R objects to, 5-57
 - running from R, 5-13, 5-16
- MapReduce programming examples in R, 5-6
- MapReduce programs
 - data compression in R, 5-61
 - data decompression in R, 5-63
- maxLoadFactor property, 3-12
- mkdir command
 - running in R, 5-28
- model.plot.R example, 5-8
- model.prep.R example, 5-9
- mv command
 - running from R, 5-29

O

- OCI Direct Path
 - See Oracle OCI Direct Path output formats
- ODI_HIVE_SESSION_JARS, for Oracle Data Integrator, 4-4
- ODI_OLH_JARS, for Oracle Data Integrator, 4-5
- ODI_OLH_SHAREDLIBS, for Oracle Data Integrator, 4-5
- ojdbc6.jar file, 2-4
- OLH_HOME environment variable, 1-6, 4-5
- OLH_OUTPUT_MODE option, 4-10
- operating system user permissions, 1-5
- Oracle, 3-20
- Oracle Data Integrator agent
 - configuring, 4-4
 - drivers, 4-4
- Oracle Data Integrator Application Adapter
 - description, 4-1
- Oracle Data Integrator Application Adapter for Hadoop
 - creating models, 4-6
 - data sources, 4-3
 - flex fields, 4-7
 - installing, 1-6
 - loading options, 4-8
 - loading output using Oracle Direct Connector for HDFS, 2-1
 - security, 4-2
 - setting up projects, 4-6
 - topology setup, 4-3
- Oracle Data Integrator Companion CD, 1-7
- Oracle Data Integrator Studio configuration, 4-6
- Oracle Data Pump
 - See Data Pump files; Data Pump format
- Oracle Database
 - connecting from R, 5-4
 - copying data from HDFS, 5-31
 - copying data to HDFS, 5-32
 - querying connections from R, 5-53
 - R access, 5-2
 - reconnecting from R, 5-62
 - user privileges, 1-5
- Oracle Database client installation, 4-5
- Oracle Database connections
 - disconnecting from R, 5-54
 - from R, 5-46, 5-48
 - reconnecting from R, 5-62
 - See also Oracle Database
- Oracle Direct Connector for HDFS
 - accessing Data Pump files, 2-2
 - configuration settings, 2-5, 2-6
 - creating external tables for, 2-2
 - Data Pump files, 2-5
 - database directories, 2-2
 - database privileges, 1-5
 - decompressing files, 2-8
 - description, 2-1
 - ExternalTable command, 2-4
 - installation, 1-3
 - loading options, 2-1
 - pattern-matching characters, 2-8
 - publishing data paths, 2-4
 - publishing Data Pump files, 2-8
 - query optimization, 2-9
 - URL connections, 2-7
- Oracle Instant Client libraries, 1-6
- Oracle Loader for Hadoop
 - description, 3-1
 - input formats, 3-2
 - installing, 1-6
 - supported database versions, 1-6
 - using in Oracle Data Integrator, 4-2, 4-5, 4-10
- Oracle Net Services, 2-7
- Oracle OCI Direct Path interface, 3-20
- Oracle OCI Direct Path output format, 3-1, 3-8, 3-16
- Oracle OCI Direct Path output formats, 3-17
- Oracle permissions, 1-5
- Oracle R Connector for Hadoop
 - access to HDFS files, 5-18
 - access to Oracle Database, 5-2
 - alphabetical list of functions, 5-3
 - categorical list of functions, 5-3
 - connecting to Oracle Database, 5-4
 - copying HDFS data, 5-4
 - debugging functions, 5-5
 - description, 5-1
 - executing scripts, 5-5
 - HDFS commands issued from, 5-4
 - installation, 1-7
 - obtaining the ORCH version number, 5-64
 - programming examples, 5-6
 - See also individual function names
- Oracle R Enterprise, 5-32, 5-34
- Oracle Software Delivery Cloud, 1-3
- Oracle Technology Network
 - certifications, 1-7

- downloads, 1-3, 1-8
- Oracle wallet, 2-4, 2-7
- ORACLE_DATAPUMP driver, 3-2, 3-10, 3-11
- ORACLE_LOADER driver, 2-1, 3-2
- oracle.hadoop.hdfs.exttab.* property
 - descriptions, 2-7
- oraclepki.jar file, 2-4, 3-6
- orahdfs-version/bin directory, 1-5
- orahdfs-version.zip file, 1-4
- OraLoader, 3-6, 3-14, 3-23
- oraloader-conf.xml document, 3-23
- OraLoaderMetadata utility program, 3-5
- oraloader-version directory, 1-6
- oraloader-version.zip file, 1-6
- ORCH package
 - installation, 1-8, 1-9
 - programming examples, 5-6
 - version numbers, 5-64
 - See also* Oracle R Connector for Hadoop
- orch.connect function
 - syntax description, 5-46
- orch.dbcon function
 - syntax description, 5-48
- orch.dbg.off function
 - syntax description, 5-50
- orch.dbg.on function
 - example, 5-6
 - syntax description, 5-51
- orch.dbg.output function
 - example, 5-6
 - syntax description, 5-52
- orch.dbinfo function
 - syntax description, 5-53
- orch.disconnect function
 - syntax description, 5-54
- orch.dryrun function
 - syntax description, 5-56
- orch.export function
 - example, 5-7, 5-8, 5-9, 5-10, 5-11
 - syntax description, 5-57
- orch.keyval function
 - example, 5-6, 5-7, 5-9, 5-10, 5-11
 - syntax description, 5-58
- orch.keyvals function
 - example, 5-9, 5-10
 - syntax description, 5-59
- orch.pack function
 - example, 5-7, 5-8, 5-9, 5-11, 5-12
 - syntax description, 5-61
- orch.reconnect function
 - syntax description, 5-62
- orch.tgz package, 1-9
- orch.unpack function
 - example, 5-7, 5-8, 5-9, 5-11, 5-12
 - syntax description, 5-63
- orch.version function
 - syntax description, 5-64
- orch.which function
 - See* orch.dbcon function
- ore.frame objects, 5-32, 5-34

- See* data frames
- ore.pull function, 5-6
- OVERRIDE_INPUTFORMAT option, 4-10
- OVERRIDE_ROW_FORMAT option, 4-8

P

- parallel processing, 1-2, 2-3, 2-9, 5-5, 5-6, 5-7, 5-30
- partitioning, 3-1, 3-22
- path, changing in R, 5-20
- pattern-matching characters in Oracle Direct Connector for HDFS, 2-8
- pca.R demo, 5-11
- pearson.R demo, 5-11
- POST_TRANSFORM_DISTRIBUTE option, 4-10
- POST_TRANSFORM_SORT option, 4-10
- PQ_DISTRIBUTE hint, 2-9
- PRE_TRANSFORM_DISTRIBUTE option, 4-10
- PRE_TRANSFORM_SORT option, 4-10
- PREPROCESSOR clause in CREATE TABLE command, 2-2
- privileges, Oracle Database, 1-5
- program execution, local dry run in R, 5-56
- projects
 - setting up in Oracle Data Integrator, 4-6
- pwd command
 - running from R, 5-36

Q

- query optimization for Oracle Direct Connector for HDFS, 2-9

R

- R connector
 - See* Oracle R Connector for Hadoop
- R Distribution, 1-9, 1-10
- R distribution, 1-8, 1-10
- R functions
 - alphabetical listing, 5-3
 - categorical listing, 5-3
 - See also* Oracle R Connector for Hadoop
- R packages
 - bitops, 1-8, 5-61, 5-63
- R programs
 - examples, 5-6
 - local execution, 5-56
- R_HOME environment variable
 - for Oracle R Connector, 1-9
- RECYCLE_ERRORS option, 4-9
- reducers
 - configuration settings, 5-6
 - examples, 5-6
 - generating key-value pairs for in R, 5-58, 5-59
 - running from R, 5-13, 5-16
- reverse engineering in Hive, 4-2, 4-6
- reverse-engineering Hive tables, 4-6
- reverse.log file, 4-7
- RKM Hive, 4-2, 4-6
- rlm.R example, 5-9

rm command
 issuing from R, 5-37
rmdir command
 issuing from R, 5-38
root directory
 changing from R, 5-41
 obtaining using R, 5-39

S

sampling data
 from Oracle Loader for Hadoop, 3-11
scripts
 debugging in R, 5-5
 executing in R, 5-5
 snippets, 4-2
SerDes JAR files, 4-5
snippets in Oracle Data Integrator, 4-2
software downloads, 1-3, 1-4, 1-7, 1-8, 1-9, 1-10, 5-61, 5-63
split.map.R example, 5-9
split.reduce.R example, 5-9
SQL*Loader, 3-7
SQL*Loader utility, 3-2
sqlnet.ora file, 2-6
Sqoop utility
 installing on a Hadoop client, 1-10
 installing on a Hadoop cluster, 1-8
SQOOP_HOME environment variable
 for Oracle R Connector, 1-9
STATIC_CONTROL option, 4-9
STOP_ON_FILE_NOT_FOUND option, 4-8
streaming data to the database, 2-4
sum.R example, 5-9

T

tables
 copying data from HDFS, 3-1
 copying data from HDFS to Oracle Database, 5-31
 copying data from Oracle Database to HDFS, 5-32
 loading data into Oracle Database, 3-4
 See also Hive tables
TEMP_DIR option, 4-10
teragen2.xy.R example, 5-10
teragen.matrix.R example, 5-10
teragen.xy.R example, 5-10
terasort.R example, 5-10
third-party licenses, 3-34
TNS_ADMIN environment variable, 2-6
TNSNAMES entries, 2-7
tnsnames.ora file, 2-6
TRANSFORM_SCRIPT option, 4-9
TRANSFORM_SCRIPT_NAME option, 4-9
transforming data
 in Oracle Data Integrator, 4-1, 4-9
TRUNCATE option, 4-8, 4-9, 4-10

U

uncompressed files, 2-8
Universal Resource Identifiers, 2-4
URL connections, for Oracle Direct Connector for HDFS, 2-7
USE_HIVE_STAGING_TABLE option, 4-10
USE_LOG option, 4-7
USE_ORACLE_STAGING_TABLE option, 4-10
USE_STAGING_TABLE option, 4-8
userlib directory, 4-6
UTL_FILE package, 1-5

V

validating data
 in Oracle Data Integrator, 4-1, 4-2, 4-9

W

wildcard characters
 in resource names, 4-8
 setting up data sources in ODI using, 4-3
 support in IKM File To Hive (Load Data), 4-2
 using to restrict input in Oracle Loader for Hadoop, 3-3

X

XML schema definition for Oracle Loader for Hadoop, 3-22
xml-reference directory, 1-7, 4-3

