

Oracle® Big Data Appliance

Software User's Guide

Release 3 (3.1)

E55983-01

August 2014

Provides an introduction to the Oracle Big Data Appliance software, tools, and administrative procedures.

Copyright © 2011, 2014, 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
 1 Introducing Oracle Big Data Appliance	
What Is Big Data?	1-1
High Variety	1-1
High Complexity	1-2
High Volume	1-2
High Velocity	1-2
The Oracle Big Data Solution	1-2
Software for Big Data	1-3
Software Component Overview	1-4
Acquiring Data for Analysis	1-4
Hadoop Distributed File System	1-5
Apache Hive	1-5
Oracle NoSQL Database	1-5
Organizing Big Data	1-6
MapReduce	1-6
Oracle Big Data Connectors	1-7
Oracle R Support for Big Data	1-8
Analyzing and Visualizing Big Data	1-9
 2 Administering Oracle Big Data Appliance	
Monitoring Multiple Clusters Using Oracle Enterprise Manager	2-1
Using the Enterprise Manager Web Interface	2-1
Using the Enterprise Manager Command-Line Interface	2-2
Managing CDH Operations Using Cloudera Manager	2-3
Monitoring the Status of Oracle Big Data Appliance	2-3
Performing Administrative Tasks	2-4
Managing Services With Cloudera Manager	2-4
Using Hadoop Monitoring Utilities	2-5
Monitoring MapReduce Jobs	2-5

Monitoring the Health of HDFS.....	2-6
Using Cloudera Hue to Interact With Hadoop	2-7
About the Oracle Big Data Appliance Software	2-8
Software Components	2-8
Unconfigured Software	2-10
About the CDH Software Services	2-11
Where Do the Services Run on a Single Rack?	2-11
Where Do the Services Run on a Multirack Cluster?	2-11
About MapReduce	2-12
Automatic Failover of the NameNode	2-13
Automatic Failover of the ResourceManager	2-14
Map and Reduce Resource Configuration	2-14
Effects of Hardware on Software Availability	2-15
Logical Disk Layout	2-15
Critical and Noncritical Nodes	2-16
First Namenode	2-16
Second NameNode	2-17
First ResourceManager	2-17
Second ResourceManager	2-17
Noncritical Nodes	2-18
Stopping and Starting Oracle Big Data Appliance	2-18
Prerequisites	2-18
Stopping Oracle Big Data Appliance	2-18
Starting Oracle Big Data Appliance	2-21
Switching From YARN to MapReduce 1	2-22
Security on Oracle Big Data Appliance	2-24
About Predefined Users and Groups	2-24
About User Authentication	2-25
About Fine-Grained Authorization	2-25
About On-Disk Encryption	2-25
Port Numbers Used on Oracle Big Data Appliance	2-26
About Puppet Security	2-26
Auditing Oracle Big Data Appliance	2-27
About Oracle Audit Vault and Database Firewall	2-27
Setting Up the Oracle Big Data Appliance Plug-in	2-28
Monitoring Oracle Big Data Appliance	2-28
Collecting Diagnostic Information for Oracle Customer Support	2-29

3 Supporting User Access to Oracle Big Data Appliance

About Accessing a Kerberos-Secured Cluster	3-1
Providing Remote Client Access to CDH	3-2
Prerequisites	3-2
Installing CDH on Oracle Exadata Database Machine	3-3
Installing a CDH Client on Any Supported Operating System	3-4
Configuring a CDH Client for an Unsecured Cluster	3-4
Configuring a CDH Client for a Kerberos-Secured Cluster	3-5
Verifying Access to a Cluster from the CDH Client	3-6

Providing Remote Client Access to Hive	3-8
Managing User Accounts	3-9
Creating Hadoop Cluster Users	3-9
Providing User Login Privileges (Optional)	3-11
Recovering Deleted Files	3-11
Restoring Files from the Trash	3-11
Changing the Trash Interval.....	3-12
Disabling the Trash Facility	3-13

4 Configuring Oracle Exadata Database Machine for Use with Oracle Big Data Appliance

About Optimizing Communications	4-1
About Applications that Pull Data Into Oracle Exadata Database Machine.....	4-1
About Applications that Push Data Into Oracle Exadata Database Machine	4-2
Prerequisites for Optimizing Communications	4-2
Specifying the InfiniBand Connections to Oracle Big Data Appliance	4-2
Specifying the InfiniBand Connections to Oracle Exadata Database Machine	4-3
Enabling SDP on Exadata Database Nodes	4-4
Configuring a JDBC Client for SDP	4-5
Creating an SDP Listener on the InfiniBand Network	4-5

5 Optimizing MapReduce Jobs Using Perfect Balance

What is Perfect Balance?	5-1
About Balancing Jobs Across Map and Reduce Tasks	5-2
Ways to Use Perfect Balance Features.....	5-2
Perfect Balance Components	5-2
Application Requirements	5-3
Getting Started with Perfect Balance	5-3
Analyzing a Job's Reducer Load	5-4
About Job Analyzer	5-4
Running Job Analyzer as a Standalone Utility	5-4
Running Job Analyzer Using Perfect Balance	5-6
Reading the Job Analyzer Report	5-8
About Configuring Perfect Balance	5-9
Running a Balanced MapReduce Job Using Perfect Balance	5-10
About Perfect Balance Reports	5-12
About Chopping	5-13
Selecting a Chopping Method	5-13
How Chopping Impacts Applications	5-14
Troubleshooting Jobs Running with Perfect Balance	5-14
Java GC Overhead Limit Exceeded Error.....	5-14
Java Out of Heap Space Errors.....	5-14
Using the Perfect Balance API	5-15
Modifying Your Java Code to Use Perfect Balance	5-15
Running Your Modified Java Code with Perfect Balance	5-16
About the Perfect Balance Examples	5-17

About the Examples in This Chapter	5-17
Extracting the Example Data Set.....	5-18
Perfect Balance Configuration Property Reference	5-18

Glossary

Index

Preface

The *Oracle Big Data Appliance Software User's Guide* describes how to manage and use the installed software.

Audience

This guide is intended for users of Oracle Big Data Appliance including:

- Application developers
- Data analysts
- Data scientists
- Database administrators
- System administrators

The *Oracle Big Data Appliance Software User's Guide* introduces Oracle Big Data Appliance installed software, features, concepts, and terminology. However, you must acquire the necessary information about administering Hadoop clusters and writing MapReduce programs from other sources.

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 Big Data Appliance Perfect Balance Java API Reference*
- *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Oracle Big Data Appliance*
- *Oracle Big Data Appliance Owner's Guide*
- *Oracle Big Data Connectors User's Guide*

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.
# prompt	The pound (#) prompt indicates a command that is run as the Linux root user.

Introducing Oracle Big Data Appliance

This chapter presents an overview of Oracle Big Data Appliance and describes the software installed on the system. This chapter contains the following sections:

- [What Is Big Data?](#)
- [The Oracle Big Data Solution](#)
- [Software for Big Data](#)
- [Acquiring Data for Analysis](#)
- [Organizing Big Data](#)
- [Analyzing and Visualizing Big Data](#)

What Is Big Data?

Using transactional data as the source of business intelligence has been commonplace for many years. As digital technology and the World Wide Web spread into every aspect of modern life, other sources of data can make important contributions to business decision making. Many businesses are looking to these new data sources. They are finding opportunities in analyzing vast amounts of data that until recently was discarded.

Big data is characterized by:

- A variety of data sources: [High Variety](#)
- A complexity of data types: [High Complexity](#)
- A high volume of data flow: [High Volume](#)
- A high velocity of data transactions: [High Velocity](#)

These characteristics pinpoint the challenges in deriving value from big data, and the differences between big data and traditional data sources that primarily provide highly structured, transactional data.

High Variety

Big data is derived from a variety of sources, such as:

- Equipment sensors: Medical, manufacturing, transportation, and other machine sensor transmissions
- Machines: Call detail records, web logs, smart meter readings, Global Positioning System (GPS) transmissions, and trading systems records

- Social media: Data streams from social media sites such as Facebook and blogging sites such as Twitter

Analysts can mine this data repeatedly as they devise new ways of extracting meaningful insights. What seems irrelevant today might prove to be highly pertinent to your business tomorrow.

Challenge: Delivering flexible systems to handle this high variety

High Complexity

As the variety of data types increases, the complexity of the system increases. The complexity of data types also increases in big data because of its low structure.

Challenge: Finding solutions that apply across a broad range of data types.

High Volume

Social media can generate terabytes of daily data. Equipment sensors and other machines can generate that much data in less than an hour.

Even traditional data sources for data warehouses, such as customer profiles from customer relationship management (CRM) systems, transactional enterprise resource planning (ERP) data, store transactions, and general ledger data, have increased tenfold in volume over the past decade.

Challenge: Providing scalability and ease in growing the system

High Velocity

Huge numbers of sensors, web logs, and other machine sources generate data continuously and at a much higher speed than traditional sources, such as individuals entering orders into a transactional database.

Challenge: Handling the data at high speed without stressing the structured systems

The Oracle Big Data Solution

Oracle Big Data Appliance is an engineered system comprising both hardware and software components. The hardware is optimized to run the enhanced big data software components.

Oracle Big Data Appliance delivers:

- A complete and optimized solution for big data
- Single-vendor support for both hardware and software
- An easy-to-deploy solution
- Tight integration with Oracle Database and Oracle Exadata Database Machine

Oracle provides a big data platform that captures, organizes, and supports deep analytics on extremely large, complex data streams flowing into your enterprise from many data sources. You can choose the best storage and processing location for your data depending on its structure, workload characteristics, and end-user requirements.


Oracle Database enables all data to be accessed and analyzed by a large user community using identical methods. By adding Oracle Big Data Appliance in front of Oracle Database, you can bring new sources of information to an existing data warehouse. Oracle Big Data Appliance is the platform for acquiring and organizing

big data so that the relevant portions with true business value can be analyzed in Oracle Database.

For maximum speed and efficiency, Oracle Big Data Appliance can be connected to Oracle Exadata Database Machine running Oracle Database. Oracle Exadata Database Machine provides outstanding performance in hosting data warehouses and transaction processing databases. Moreover, Oracle Exadata Database Machine can be connected to Oracle Exalytics In-Memory Machine for the best performance of business intelligence and planning applications. The InfiniBand connections between these engineered systems provide high parallelism, which enables high-speed data transfer for batch or query workloads.

Figure 1–1 shows the relationships among these engineered systems.

Figure 1–1 Oracle Engineered Systems for Big Data



Software for Big Data

The **Oracle Linux** operating system and Cloudera's Distribution including Apache Hadoop (CDH) underlie all other software components installed on Oracle Big Data Appliance. **CDH** is an integrated stack of components that have been tested and packaged to work together.

CDH has a batch processing infrastructure that can store files and distribute work across a set of computers. Data is processed on the same computer where it is stored. In a single Oracle Big Data Appliance rack, CDH distributes the files and workload across 18 servers, which compose a **cluster**. Each server is a node in the cluster.

The software framework consists of these primary components:

- **File system:** The **Hadoop Distributed File System (HDFS)** is a highly scalable file system that stores large files across multiple servers. It achieves reliability by

replicating data across multiple servers without RAID technology. It runs on top of the Linux file system on Oracle Big Data Appliance.

- **MapReduce engine:** The **MapReduce** engine provides a platform for the massively parallel execution of algorithms written in Java. Oracle Big Data Appliance 3.0 runs **YARN** by default.
- **Administrative framework:** **Cloudera Manager** is a comprehensive administrative tool for CDH. In addition, you can use Oracle Enterprise Manager to monitor both the hardware and software on Oracle Big Data Appliance.
- **Apache projects:** CDH includes Apache projects for MapReduce and HDFS, such as **Hive**, **Pig**, **Oozie**, **ZooKeeper**, **HBase**, **Sqoop**, and **Spark**.
- **Cloudera applications:** Oracle Big Data Appliance installs all products included in Cloudera Enterprise Data Hub Edition, including **Impala**, **Search**, and **Navigator**.

CDH is written in Java, and Java is the language for applications development. However, several CDH utilities and other software available on Oracle Big Data Appliance provide graphical, web-based, and other language interfaces for ease of use.

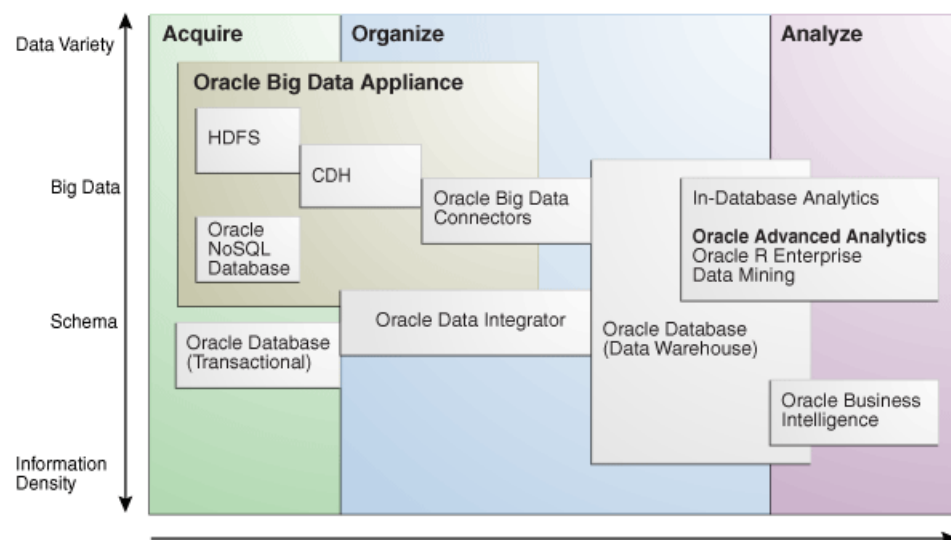
Software Component Overview

The major software components perform three basic tasks:

- Acquire
- Organize
- Analyze and visualize

The best tool for each task depends on the density of the information and the degree of structure. [Figure 1–2](#) shows the relationships among the tools and identifies the tasks that they perform.

Figure 1–2 Oracle Big Data Appliance Software Overview



Acquiring Data for Analysis

Databases used for online transaction processing (OLTP) are the traditional data sources for data warehouses. The Oracle solution enables you to analyze traditional data stores with big data in the same Oracle data warehouse. Relational data continues

to be an important source of business intelligence, although it runs on separate hardware from Oracle Big Data Appliance.

Oracle Big Data Appliance provides these facilities for capturing and storing big data:

- [Hadoop Distributed File System](#)
- [Apache Hive](#)
- [Oracle NoSQL Database](#)

Hadoop Distributed File System

Cloudera's Distribution including Apache Hadoop (CDH) on Oracle Big Data Appliance uses the Hadoop Distributed File System (HDFS). HDFS stores extremely large files containing record-oriented data. On Oracle Big Data Appliance, HDFS splits large data files into chunks of 256 megabytes (MB), and replicates each chunk across three different nodes in the cluster. The size of the chunks and the number of replications are configurable.

Chunking enables HDFS to store files that are larger than the physical storage of one server. It also allows the data to be processed in parallel across multiple computers with multiple processors, all working on data that is stored locally. Replication ensures the high availability of the data: if a server fails, the other servers automatically take over its work load.

HDFS is typically used to store all types of big data.

See Also:

- For conceptual information about Hadoop technologies, refer to this third-party publication:
Hadoop: The Definitive Guide, Third Edition by Tom White (O'Reilly Media Inc., 2012, ISBN: 978-1449311520).
- For documentation about Cloudera's Distribution including Apache Hadoop, see the Cloudera library at <http://oracle.cloudera.com/>

Apache Hive

Hive is an open-source data warehouse that supports data summarization, ad hoc querying, and data analysis of data stored in HDFS. It uses a SQL-like language called [HiveQL](#). An interpreter generates MapReduce code from the HiveQL queries. By storing data in Hive, you can avoid writing MapReduce programs in Java.

Hive is a component of CDH and is always installed on Oracle Big Data Appliance. Oracle Big Data Connectors can access Hive tables.

Oracle NoSQL Database

Oracle NoSQL Database is a distributed key-value database built on the proven storage technology of Berkeley DB Java Edition. Whereas HDFS stores unstructured data in very large files, Oracle NoSQL Database indexes the data and supports transactions. But unlike Oracle Database, which stores highly structured data, Oracle NoSQL Database has relaxed consistency rules, no schema structure, and only modest support for joins, particularly across storage nodes.

NoSQL databases, or "Not Only SQL" databases, have developed over the past decade specifically for storing big data. However, they vary widely in implementation. Oracle NoSQL Database has these characteristics:

- Uses a system-defined, consistent hash index for data distribution
- Supports high availability through replication
- Provides single-record, single-operation transactions with relaxed consistency guarantees
- Provides a Java API

Oracle NoSQL Database is designed to provide highly reliable, scalable, predictable, and available data storage. The key-value pairs are stored in shards or partitions (that is, subsets of data) based on a primary key. Data on each shard is replicated across multiple storage nodes to ensure high availability. Oracle NoSQL Database supports fast querying of the data, typically by key lookup.

An intelligent driver links the NoSQL database with client applications and provides access to the requested key-value on the storage node with the lowest latency.

Oracle NoSQL Database includes hashing and balancing algorithms to ensure proper data distribution and optimal load balancing, replication management components to handle storage node failure and recovery, and an easy-to-use administrative interface to monitor the state of the database.

Oracle NoSQL Database is typically used to store customer profiles and similar data for identifying and analyzing big data. For example, you might log in to a website and see advertisements based on your stored customer profile (a record in Oracle NoSQL Database) and your recent activity on the site (web logs currently streaming into HDFS).

Oracle NoSQL Database is an optional component of Oracle Big Data Appliance and runs on a separate cluster from CDH.

See Also:

- *Oracle NoSQL Database Getting Started Guide* at
<http://docs.oracle.com/cd/NOSQL/html/index.html>
- *Oracle Big Data Appliance Licensing Information*

Organizing Big Data

Oracle Big Data Appliance provides several ways of organizing, transforming, and reducing big data for analysis:

- [MapReduce](#)
- [Oracle Big Data Connectors](#)
- [Oracle R Support for Big Data](#)

MapReduce

The MapReduce engine provides a platform for the massively parallel execution of algorithms written in Java. MapReduce uses a parallel programming model for processing data on a distributed system. It can process vast amounts of data quickly and can scale linearly. It is particularly effective as a mechanism for batch processing of unstructured and semistructured data. MapReduce abstracts lower-level operations into computations over a set of keys and values.

Although big data is often described as unstructured, incoming data always has some structure. However, it does not have a fixed, predefined structure when written to HDFS. Instead, MapReduce creates the desired structure as it reads the data for a particular job. The same data can have many different structures imposed by different MapReduce jobs.

A simplified description 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 is produced as key-value pairs. The Map phase is logically run in parallel over each record, whereas the Reduce phase is run in parallel over all key values.

Note: Oracle Big Data Appliance uses the Yet Another Resource Negotiator (YARN) implementation of MapReduce by default. You have the option of using classic MapReduce (MRv1) instead. You cannot use both implementations in the same cluster; you can activate either the MapReduce or the YARN service.

Oracle Big Data Connectors

Oracle Big Data Connectors facilitate data access between data stored in CDH and Oracle Database. The connectors are licensed separately from Oracle Big Data Appliance and include:

- [Oracle SQL Connector for Hadoop Distributed File System](#)
- [Oracle Loader for Hadoop](#)
- [Oracle XQuery for Hadoop](#)
- [Oracle R Advanced Analytics for Hadoop](#)
- [Oracle Data Integrator Application Adapter for Hadoop](#)

See Also: *Oracle Big Data Connectors User's Guide*

Oracle SQL Connector for Hadoop Distributed File System

Oracle SQL Connector for Hadoop Distributed File System (Oracle SQL Connector for HDFS) provides read access to HDFS from an Oracle database using **external tables**.

An external table is an Oracle Database object that identifies the location of data outside of the 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 the database. External tables are often used to stage data to be transformed during a database load.

You can use Oracle SQL Connector for HDFS to:

- Access data stored in HDFS files
- Access Hive tables.
- Access comma-separated value (CSV) files generated by Oracle Loader for Hadoop

- Load data extracted and transformed by Oracle Data Integrator

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 partitions the data and transforms it into a database-ready format in Hadoop. It optionally sorts records by primary key before loading the data or creating output files.

You can use Oracle Loader for Hadoop as either a Java program or a command-line utility. The load runs as a MapReduce job on the Hadoop cluster.

Oracle Loader for Hadoop also reads from and writes to Oracle Data Pump files.

Oracle Data Integrator Application Adapter for Hadoop

Oracle Data Integrator (ODI) extracts, transforms, and loads data into Oracle Database from a wide range of sources.

In ODI, a knowledge module (KM) is a code template dedicated to a specific task in the data integration process. You use Oracle Data Integrator Studio to load, select, and configure the KMs for your particular application. More than 150 KMs are available to help you acquire data from a wide range of third-party databases and other data repositories. You only need to load a few KMs for any particular job.

Oracle Data Integrator Application Adapter for Hadoop contains the KMs specifically for use with big data.

Oracle XQuery for Hadoop

Oracle XQuery for Hadoop runs transformations expressed in the XQuery language by translating them into a series of MapReduce jobs, which are executed in parallel on the Hadoop cluster. The input data can be located in HDFS or Oracle NoSQL Database. Oracle XQuery for Hadoop can write the transformation results to HDFS, Oracle NoSQL Database, or Oracle Database.

Oracle R Advanced Analytics for Hadoop

Oracle R Advanced Analytics for Hadoop is a collection of R packages that provides:

- Interfaces to work with Hive tables, Apache Hadoop compute infrastructure, local R environment and database tables
- Predictive analytic techniques written in R or Java as Hadoop MapReduce jobs that can be applied to data in HDFS files

Using simple R functions, you can copy data between R memory, the local file system, HDFS, and Hive. You can write mappers and reducers in R, schedule these R programs to execute as Hadoop MapReduce jobs, and return the results to any of those locations.

Oracle R Support for Big Data

R is an open-source language and environment for statistical analysis and graphing. It provides linear and nonlinear modeling, standard statistical methods, time-series analysis, classification, clustering, and graphical data displays. Thousands of open-source packages are available in the Comprehensive R Archive Network (CRAN) for a spectrum of applications, such as bioinformatics, spatial statistics, and financial and marketing analysis. The popularity of R has increased as its functionality matured to rival that of costly proprietary statistical packages.

Analysts typically use R on a PC, which limits the amount of data and the processing power available for analysis. Oracle eliminates this restriction by extending the R platform to directly leverage Oracle Big Data Appliance. Oracle R Distribution is installed on all nodes of Oracle Big Data Appliance.

Oracle R Advanced Analytics for Hadoop provides R users with high-performance, native access to HDFS and the MapReduce programming framework, which enables R programs to run as MapReduce jobs on vast amounts of data. Oracle R Advanced Analytics for Hadoop is included in the Oracle Big Data Connectors. See "[Oracle R Advanced Analytics for Hadoop](#)" on page 1-8.

Oracle R Enterprise is a component of the Oracle Advanced Analytics option to Oracle Database. It provides:

- Transparent access to database data for data preparation and statistical analysis from R
- Execution of R scripts at the database server, accessible from both R and SQL
- A wide range of predictive and data mining in-database algorithms

Oracle R Enterprise enables you to store the results of your analysis of big data in an Oracle database, or accessed for display in dashboards and applications.

Both Oracle R Advanced Analytics for Hadoop and Oracle R Enterprise make Oracle Database and the Hadoop computational infrastructure available to statistical users without requiring them to learn the native programming languages of either one.

See Also:

- For information about R, go to
<http://www.r-project.org/>
- For information about Oracle R Enterprise, go to
http://docs.oracle.com/cd/E40980_01/welcome.html

Analyzing and Visualizing Big Data

After big data is transformed and loaded in Oracle Database, you can use the full spectrum of Oracle business intelligence solutions and decision support products to further analyze and visualize all your data.

See Also:

- Oracle Business Intelligence website at
<http://www.oracle.com/us/solutions/ent-performance-bi/business-intelligence/index.html>
- Data Warehousing and Business Intelligence in the Oracle Database Documentation Library at
http://www.oracle.com/pls/db112/portal.portal_db?selected=6&frame=

Administering Oracle Big Data Appliance

This chapter provides information about the software and services installed on Oracle Big Data Appliance. It contains these sections:

- [Monitoring Multiple Clusters Using Oracle Enterprise Manager](#)
- [Managing CDH Operations Using Cloudera Manager](#)
- [Using Hadoop Monitoring Utilities](#)
- [Using Cloudera Hue to Interact With Hadoop](#)
- [About the Oracle Big Data Appliance Software](#)
- [About the CDH Software Services](#)
- [Effects of Hardware on Software Availability](#)
- [Stopping and Starting Oracle Big Data Appliance](#)
- [Switching From YARN to MapReduce 1](#)
- [Security on Oracle Big Data Appliance](#)
- [Auditing Oracle Big Data Appliance](#)
- [Collecting Diagnostic Information for Oracle Customer Support](#)

Monitoring Multiple Clusters Using Oracle Enterprise Manager

An Oracle Enterprise Manager plug-in enables you to use the same system monitoring tool for Oracle Big Data Appliance as you use for Oracle Exadata Database Machine or any other Oracle Database installation. With the plug-in, you can view the status of the installed software components in tabular or graphic presentations, and start and stop these software services. You can also monitor the health of the network and the rack components.

Oracle Enterprise Manager enables you to monitor all Oracle Big Data Appliance racks on the same InfiniBand fabric. It provides summary views of both the rack hardware and the software layout of the logical clusters.

Using the Enterprise Manager Web Interface

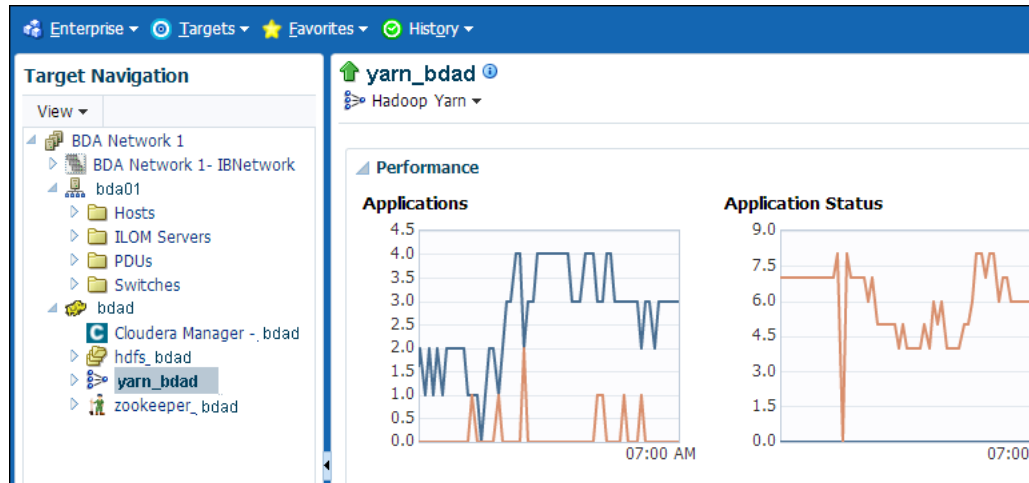
After opening Oracle Enterprise Manager web interface, logging in, and selecting a target cluster, you can drill down into these primary areas:

- **InfiniBand network:** Network topology and status for InfiniBand switches and ports. See [Figure 2-1](#).
- **Hadoop cluster:** Software services for HDFS, MapReduce, and ZooKeeper.

- **Oracle Big Data Appliance rack:** Hardware status including server hosts, Oracle Integrated Lights Out Manager (Oracle ILOM) servers, power distribution units (PDUs), and the Ethernet switch.

Figure 2–1 shows a small section of the cluster home page.

Figure 2–1 YARN Page in Oracle Enterprise Manager



To monitor Oracle Big Data Appliance using Oracle Enterprise Manager:

1. Download and install the plug-in. See *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Oracle Big Data Appliance*.
2. Log in to Oracle Enterprise Manager as a privileged user.
3. From the Targets menu, choose **Big Data Appliance** to view the Big Data page. You can see the overall status of the targets already discovered by Oracle Enterprise Manager.
4. Select a target cluster to view its detail pages.
5. Expand the target navigation tree to display the components. Information is available at all levels.
6. Select a component in the tree to display its home page.
7. To change the display, choose an item from the drop-down menu at the top left of the main display area.

See Also: *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Oracle Big Data Appliance* for installation instructions and use cases.

Using the Enterprise Manager Command-Line Interface

The Enterprise Manager command-line interface (`emcli`) is installed on Oracle Big Data Appliance along with all the other software. It provides the same functionality as the web interface. You must provide credentials to connect to Oracle Management Server.

To get help, enter `emcli help`.

See Also: *Oracle Enterprise Manager Command Line Interface Guide*

Managing CDH Operations Using Cloudera Manager

Cloudera Manager is installed on Oracle Big Data Appliance to help you with Cloudera's Distribution including Apache Hadoop (CDH) operations. Cloudera Manager provides a single administrative interface to all Oracle Big Data Appliance servers configured as part of the Hadoop cluster.

Cloudera Manager simplifies the performance of these administrative tasks:

- Monitor jobs and services
- Start and stop services
- Manage security and Kerberos credentials
- Monitor user activity
- Monitor the health of the system
- Monitor performance metrics
- Track hardware use (disk, CPU, and RAM)

Cloudera Manager runs on the ResourceManager node (node03) and is available on port 7180.

To use Cloudera Manager:

1. Open a browser and enter a URL like the following:

`http://bda1node03.example.com:7180`

In this example, `bda1` is the name of the appliance, `node03` is the name of the server, `example.com` is the domain, and `7180` is the default port number for Cloudera Manager.

2. Log in with a user name and password for Cloudera Manager. Only a user with administrative privileges can change the settings. Other Cloudera Manager users can view the status of Oracle Big Data Appliance.

See Also: *Cloudera Manager Monitoring and Diagnostics Guide* at

<http://www.cloudera.com/content/cloudera-content/cloudera-docs/CM5/latest/Cloudera-Manager-Diagnostics-Guide/Cloudera-Manager-Diagnostics-Guide.html>

or click **Help** on the Cloudera Manager Support menu

Monitoring the Status of Oracle Big Data Appliance

In Cloudera Manager, you can choose any of the following pages from the menu bar across the top of the display:

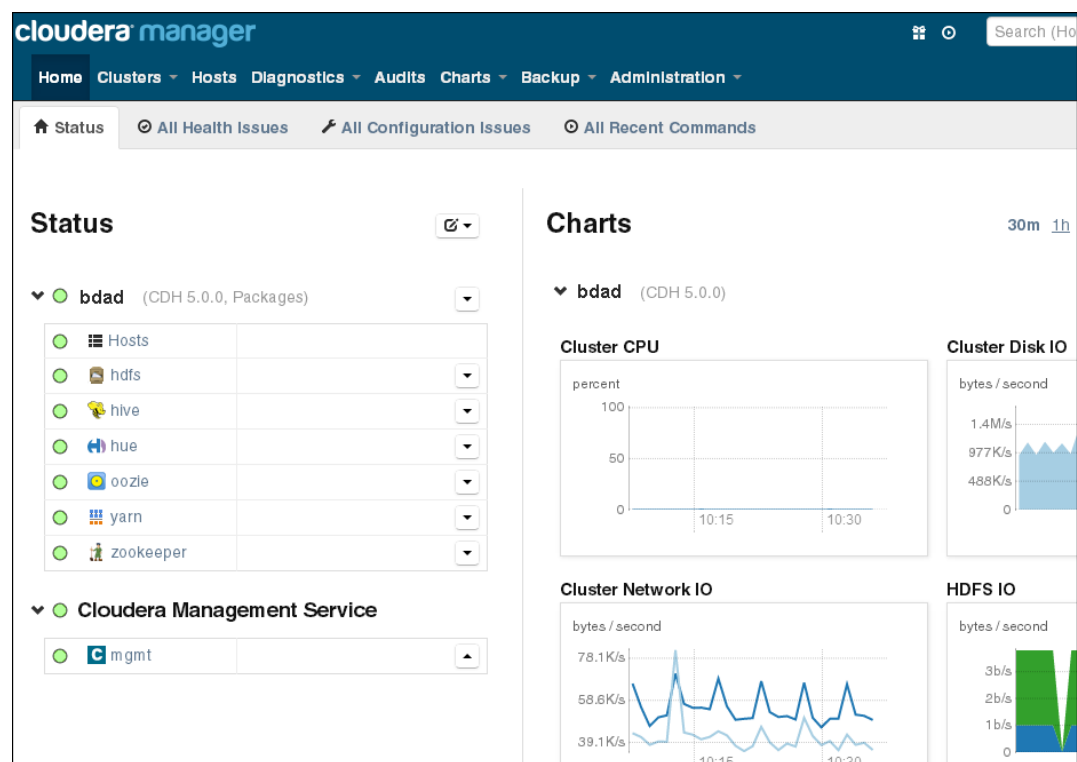
- **Home:** Provides a graphic overview of activities and links to all services controlled by Cloudera Manager. See [Figure 2-2](#).
- **Clusters:** Accesses the services on multiple clusters.
- **Hosts:** Monitors the health, disk usage, load, physical memory, swap space, and other statistics for all servers in the cluster.
- **Diagnostics:** Accesses events and logs. Cloudera Manager collects historical information about the systems and services. You can search for a particular phrase for a selected server, service, and time period. You can also select the minimum

severity level of the logged messages included in the search: TRACE, DEBUG, INFO, WARN, ERROR, or FATAL.

- **Audits:** Displays the audit history log for a selected time range. You can filter the results by user name, service, or other criteria, and download the log as a CSV file.
- **Charts:** Enables you to view metrics from the Cloudera Manager time-series data store in a variety of chart types, such as line and bar.
- **Backup:** Accesses snapshot policies and scheduled replications.
- **Administration:** Provides a variety of administrative options, including Settings, Alerts, Users, and Kerberos.

Figure 2–2 shows the Cloudera Manager home page.

Figure 2–2 Cloudera Manager Home Page



Performing Administrative Tasks

As a Cloudera Manager administrator, you can change various properties for monitoring the health and use of Oracle Big Data Appliance, add users, and set up Kerberos security.

To access Cloudera Manager Administration:

1. Log in to Cloudera Manager with administrative privileges.
2. Select the Administration tab.

Managing Services With Cloudera Manager

Cloudera Manager provides the interface for managing these services:

- HDFS

- Hive
- Hue
- Oozie
- YARN
- ZooKeeper

You can use Cloudera Manager to change the configuration of these services, stop, and restart them. Additional services are also available, which require configuration before you can use them. See ["Unconfigured Software"](#) on page 2-10.

Note: Manual edits to Linux service scripts or Hadoop configuration files do not affect these services. You must manage and configure them using Cloudera Manager.

Using Hadoop Monitoring Utilities

You also have the option of using the native Hadoop utilities. These utilities are read-only and do not require authentication.

Cloudera Manager provides an easy way to obtain the correct URLs for these utilities. On the YARN service page, expand the Web UI submenu.

Monitoring MapReduce Jobs

You can monitor MapReduce jobs using the resource manager interface.

To monitor MapReduce jobs:

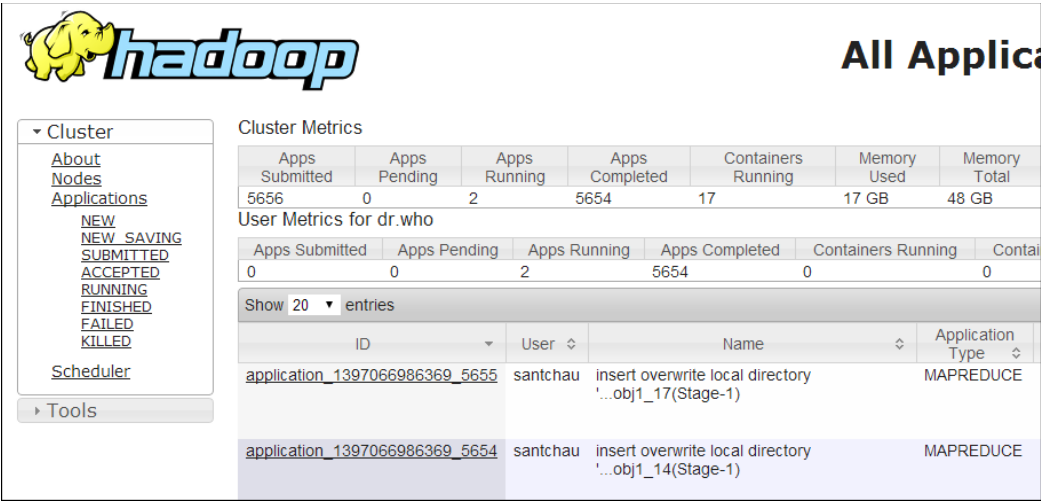
- Open a browser and enter a URL like the following:

`http://bda1node03.example.com:8088`

In this example, bda1 is the name of the rack, node03 is the name of the server where the YARN resource manager runs, and 8088 is the default port number for the user interface.

[Figure 2-3](#) shows the resource manager interface.

Figure 2–3 YARN Resource Manager Interface



Monitoring the Health of HDFS

You can monitor the health of the Hadoop file system by using the DFS health utility on the first two nodes of a cluster.

To monitor HDFS:

- Open a browser and enter a URL like the following:

http://bda1node01.example.com:50070

In this example, bda1 is the name of the rack, node01 is the name of the server where the dfshealth utility runs, and 50070 is the default port number for the user interface.

Figure 2–3 shows the DFS health utility interface.

Figure 2–4 DFS Health Utility

Hadoop

Overview

Datanodes

Snapshot

Startup Progress

Utilities

Overview '0.0.0.0:8020' (active)

Started:	Thu Apr 03 11:38:34 PDT 2014
Version:	2.3.0-cdh5.0.0, r8e266e052e423af592871e2dfe09d54c03f6a0e8
Compiled:	2014-03-28T04:30Z by jenkins from (no branch)
Cluster ID:	cluster1
Block Pool ID:	BP-1111518783-192.168.41.61-1396506058331

Summary

Security is off.

Safemode is off.

79286 files and directories, 133876 blocks = 213162 total filesystem object(s).

Using Cloudera Hue to Interact With Hadoop

Hue runs in a browser and provides an easy-to-use interface to several applications to support interaction with Hadoop and HDFS. You can use Hue to perform any of the following tasks:

- Query Hive data stores
- Create, load, and delete Hive tables
- Work with HDFS files and directories
- Create, submit, and monitor MapReduce jobs
- Monitor MapReduce jobs
- Create, edit, and submit workflows using the Oozie dashboard
- Manage users and groups

Hue is automatically installed and configured on Oracle Big Data Appliance. It runs on port 8888 of the ResourceManager node (node03).

To use Hue:

1. Log in to Cloudera Manager and click the **hue** service on the Home page.
2. On the hue page, click Hue Web UI.
3. Bookmark the Hue URL, so that you can open Hue directly in your browser. The following URL is an example:

`http://bda1node03.example.com:8888`

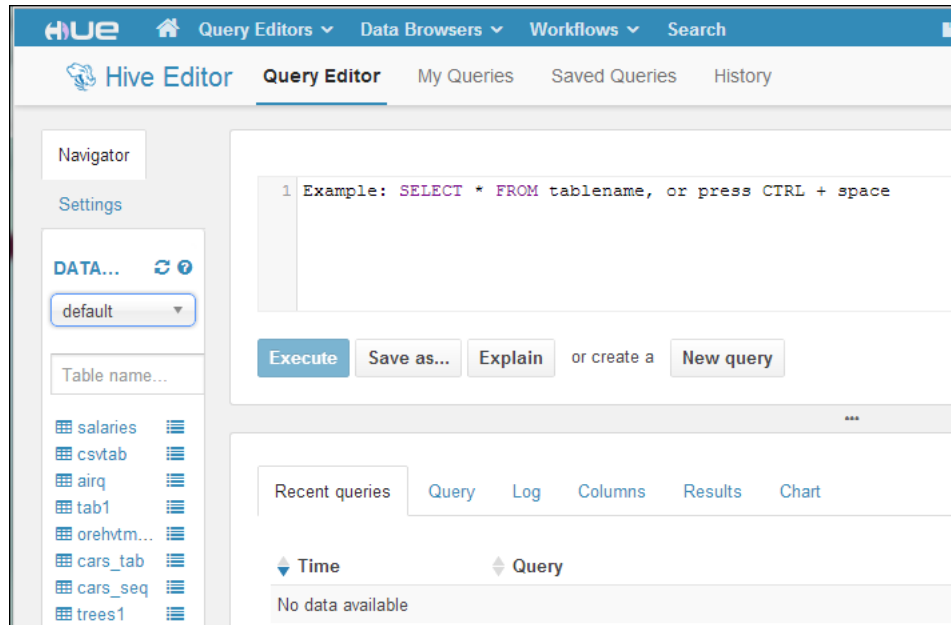
4. Log in with your Hue credentials.

Oracle Big Data Appliance is not configured initially with any Hue user accounts. The first user who connects to Hue can log in with any user name and password,

and automatically becomes an administrator. This user can create other user and administrator accounts.

Figure 2–5 shows the Hive Query Editor.

Figure 2–5 *Hive Query Editor*



See Also: *Hue User Guide* at

<http://archive-primary.cloudera.com/cdh5/cdh/5/hue/user-guide/>

About the Oracle Big Data Appliance Software

The following sections identify the software installed on Oracle Big Data Appliance. Some components operate with Oracle Database 11.2.0.2 and later releases.

This section contains the following topics:

- [Software Components](#)
- [Unconfigured Software](#)

Software Components

These software components are installed on all servers in the cluster. Oracle Linux, required drivers, firmware, and hardware verification utilities are factory installed. All other software is installed on site. The optional software components may not be configured in your installation.

Note: You do not need to install additional software on Oracle Big Data Appliance. Doing so may result in a loss of warranty and support. See the *Oracle Big Data Appliance Owner's Guide*.

Base image software:

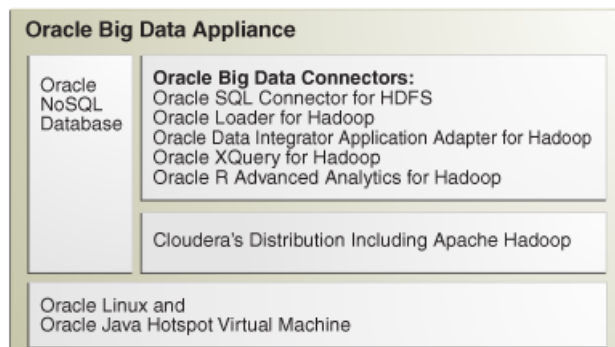
- Oracle Linux 6.4 (upgrades stay at 5.8) with Oracle Unbreakable Enterprise Kernel version 2 (UEK2)
- Java [HotSpot](#) Virtual Machine 7 version 25 (JDK 7u25)
- [Oracle R Distribution](#) 3.0.1-2
- [MySQL Database](#) 5.5.35 Advanced Edition
- Puppet, firmware, Oracle Big Data Appliance utilities
- Oracle InfiniBand software

Mammoth installation:

- Cloudera's Distribution including Apache Hadoop Release 5 (5.0) including:
 - Apache Hive 0.12
 - Apache HBase
 - [Apache Sentry](#)
 - Apache Spark
 - [Cloudera Impala](#)
 - [Cloudera Search](#) 1.2.0
- Cloudera Manager Release 5 (5.0) including:
 - [Cloudera Navigator](#)
- [Oracle Database Instant Client](#) 12.1
- Oracle NoSQL Database Community Edition or Enterprise Edition 12c Release 1 Version 3.0.5 (optional)
- Oracle Big Data Connectors 3.0 (optional):
 - Oracle SQL Connector for Hadoop Distributed File System (HDFS)
 - Oracle Loader for Hadoop
 - Oracle Data Integrator Agent 11.1.1.7.0
 - Oracle XQuery for Hadoop
 - Oracle R Advanced Analytics for Hadoop

See Also: *Oracle Big Data Appliance Owner's Guide* for information about the Mammoth utility

[Figure 2–6](#) shows the relationships among the major components.

Figure 2–6 Major Software Components of Oracle Big Data Appliance

Unconfigured Software

Your Oracle Big Data Appliance license includes all components in Cloudera Enterprise Data Hub Edition. All CDH components are installed automatically by the Mammoth utility. Do not download them from the Cloudera website.

See Also: For a complete list of CDH components:

<http://www.cloudera.com/content/cloudera/en/products-and-services/product-comparison.html>

However, you must use Cloudera Manager to add the following services before you can use them:

- [Apache Flume](#)
- [Apache HBase](#)
- [Apache Spark](#)
- [Apache Sqoop](#)
- [Cloudera Impala](#)
- [Cloudera Navigator](#)
- [Cloudera Search \(Solr\)](#)

To add a service:

1. Log in to Cloudera Manager as the admin user.
2. On the Home page, expand the cluster menu in the left panel and choose **Add a Service** to open the Add Service wizard. The first page lists the services you can add.
3. Follow the steps of the wizard.

You can find the RPM files on the first server of each cluster in `/opt/oracle/BDAMammoth/bdarepo/RPMS/noarch`.

See Also: *CDH5 Installation and Configuration Guide* for configuration procedures at

<http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH5/latest/CDH5-Installation-Guide/CDH5-Installation-Guide.html>

About the CDH Software Services

All services are installed on all nodes in a CDH cluster, but individual services run only on designated nodes. There are slight variations in the location of the services depending on the configuration of the cluster.

This section describes the services in a default YARN configuration.

This section contains the following topics:

- [Where Do the Services Run on a Single Rack?](#)
- [Where Do the Services Run on a Multirack Cluster?](#)
- [About MapReduce](#)
- [Automatic Failover of the NameNode](#)
- [Automatic Failover of the ResourceManager](#)

Where Do the Services Run on a Single Rack?

[Table 2–1](#) identifies the services in CDH clusters configured within a single rack, including starter racks and clusters with six nodes. Node01 is the first server in the cluster (server 1, 7, or 10), and nodemn is the last server in the cluster (server 6, 9, 12, or 18).

Table 2–1 Service Locations for One or More CDH Clusters in a Single Rack

Node01	Node02	Node03	Node04	Node05 to nn
Balancer		Cloudera Manager Server		
Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent
DataNode	DataNode	DataNode	DataNode	DataNode
Failover Controller	Failover Controller	JobHistory	Hive, Hue, Oozie, Solr	
JournalNode	JournalNode	JournalNode		
	MySQL Backup	MySQL Primary		
NameNode	NameNode			
NodeManager ¹	NodeManager ¹	NodeManager	NodeManager	NodeManager
			Oracle Data Integrator Agent	
Puppet	Puppet	Puppet	Puppet	Puppet
Puppet Master		ResourceManager	ResourceManager	
ZooKeeper	ZooKeeper	ZooKeeper		

¹ Starter racks and six-node clusters only, with reduced allocated resources

Where Do the Services Run on a Multirack Cluster?

When multiple racks are configured as a single CDH cluster, some critical services are installed on the second rack.

[Table 2–2](#) identifies the location of services on the first rack of a multirack cluster.

Table 2–2 Service Locations in the First Rack of a Multirack CDH Cluster

Node01	Node02	Node03	Node04	Node05 to nn ¹
Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent
	Cloudera Manager Server			
DataNode	DataNode	DataNode	DataNode	DataNode
Failover Controller				
JournalNode	JournalNode			
NameNode	MySQL Primary			
NodeManager	NodeManager	NodeManager	NodeManager	NodeManager
Puppet	Puppet	Puppet	Puppet	Puppet
Puppet Master	ResourceManager			
ZooKeeper	ZooKeeper			

¹ *nn* includes the servers in additional racks.

Table 2–3 shows the service locations in the second rack of a multirack cluster.

Table 2–3 Service Locations in the Second Rack of a Multirack CDH Cluster

Node01	Node02	Node03	Node04	Node05 to nn
Balancer				
Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent	Cloudera Manager Agent
DataNode	DataNode	DataNode	DataNode	DataNode
Failover Controller				
JournalNode	Hive, Hue, Oozie, Solr			
MySQL Backup				
NameNode				
NodeManager ¹	NodeManager ¹	NodeManager	NodeManager	NodeManager
	Oracle Data Integrator Agent			
Puppet	Puppet	Puppet	Puppet	Puppet
Puppet Master	ResourceManager			
ZooKeeper				

¹ Starter racks and six-node clusters only, with reduced allocated resources

About MapReduce

Yet Another Resource Negotiator (YARN) is the default version of MapReduce, starting with Oracle Big Data Appliance 3.0. Oracle recommends YARN for production systems. MapReduce 1 (MRv1) is supported for backward compatibility, but you must install it manually. Do not run both versions of MapReduce (YARN and MRv1) on the same cluster.

MapReduce applications developed using MapReduce 1 may require recompilation to run under YARN.

The ResourceManager performs all resource management tasks. An MRAppMaster performs the job management tasks. Each job has its own MRAppMaster. The NodeManager has containers that can run a map task, a reduce task, or an MRAppMaster. The NodeManager can dynamically allocate containers using the available memory. This architecture results in improved scalability and better use of the cluster than MRv1.

YARN also manages resources for Spark and Impala.

MRv1 manages MapReduce jobs exclusively. The JobTracker combines the functions of the ResourceManager and MRAppMaster. TaskTrackers perform the function of the NodeManagers.

See Also: "Running Existing Applications on Hadoop 2 YARN" at

<http://hortonworks.com/blog/running-existing-applications-on-hadoop-2-yarn/>

Automatic Failover of the NameNode

The NameNode is the most critical process because it keeps track of the location of all data. Without a healthy NameNode, the entire cluster fails. Apache Hadoop v0.20.2 and earlier are vulnerable to failure because they have a single name node.

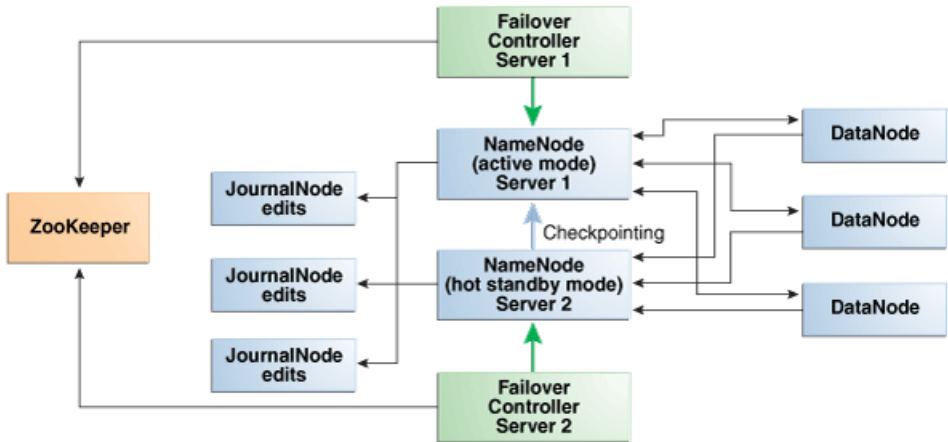
Cloudera's Distribution including Apache Hadoop Version 4 (CDH5) reduces this vulnerability by maintaining redundant NameNodes. The data is replicated during normal operation as follows:

- CDH maintains redundant NameNodes on the first two nodes of a cluster. One of the NameNodes is in active mode, and the other NameNode is in hot standby mode. If the active NameNode fails, then the role of active NameNode automatically fails over to the standby NameNode.
- The NameNode data is written to a mirrored partition so that the loss of a single disk can be tolerated. This mirroring is done at the factory as part of the operating system installation.
- The active NameNode records all changes to the file system metadata in at least two JournalNode processes, which the standby NameNode reads. There are three JournalNodes, which run on the first three nodes of each cluster.
- The changes recorded in the journals are periodically consolidated into a single fsimage file in a process called **checkpointing**.

Note: Oracle Big Data Appliance 2.0 and later releases do not support the use of an external NFS filer for backups and do not use NameNode federation.

Figure 2–7 shows the relationships among the processes that support automatic failover of the NameNode.

Figure 2–7 Automatic Failover of the NameNode on Oracle Big Data Appliance




Automatic Failover of the ResourceManager

The ResourceManager allocates resources for application tasks and application masters across the cluster. Like the NameNode, the ResourceManager is a critical point of failure for the cluster. If all ResourceManagers fail, then all jobs stop running. Oracle Big Data Appliance 3.0 supports ResourceManager High Availability in Cloudera 5 to reduce this vulnerability.

CDH maintains redundant ResourceManager services on node03 and node04. One of the services is in active mode, and the other service is in hot standby mode. If the active service fails, then the role of active ResourceManager automatically fails over to the standby service. No failover controllers are required.

Figure 2–7 shows the relationships among the processes that support automatic failover of the ResourceManager.

Figure 2–8 Automatic Failover of the ResourceManager on Oracle Big Data Appliance



Map and Reduce Resource Configuration

The NodeManager services are allocated a fixed amount of memory and CPU core voltage (Vcore), as shown in Table 2–4.

Table 2–4 Resource Allocation for NodeManagers

Appliance Model	Memory	VCore
Oracle Big Data Appliance X4-2	40 GB	32V
Oracle Big Data Appliance X3-2	40 GB	32V

Table 2–4 (Cont.) Resource Allocation for NodeManagers

Appliance Model	Memory	VCore
Oracle Big Data Appliance (Sun Fire X4270 M2-based rack)	30 GB	24V

Each node in a cluster has a maximum number of map and reduce tasks that are allowed to run simultaneously. [Table 2–5](#) shows the default configuration of resources for the MapReduce service on Oracle Big Data Appliance X4-2 and X3-2.

Table 2–5 Maximum Map and Reduce Tasks on Oracle Big Data Appliance X3-2

Node	6-Node Cluster	Larger Clusters ¹
Node01 and Node02	14 map 8 reduce	None (no NodeManager)
Node03 and Node04	10 map 6 reduce	10 map 6 reduce
Node05 to Nodenn	20 map 13 reduce	20 map 13 reduce

¹ 9 or more nodes

[Table 2–6](#) shows the default configuration of resources for the MapReduce service on Oracle Big Data Appliance Sun Fire X4270 M2-based racks.

Table 2–6 Maximum Map and Reduce Tasks on Sun Fire X4270 M2-Based Racks

Node	6-Node Cluster	Larger Clusters ¹
Node01 and Node02	10 map 6 reduce	None
Node03 and Node04	7 map 4 reduce	7 map 4 reduce
Node05 to Nodenn	15 map 10 reduce	15 map 10 reduce

¹ 9 or more nodes

Effects of Hardware on Software Availability

The effects of a server failure vary depending on the server's function within the CDH cluster. Oracle Big Data Appliance servers are more robust than commodity hardware, so you should experience fewer hardware failures. This section highlights the most important services that run on the various servers of the primary rack. For a full list, see ["Where Do the Services Run on a Single Rack?"](#) on page 2-11.

Note: In a multirack cluster, some critical services run on the first server of the second rack. See ["Where Do the Services Run on a Multirack Cluster?"](#) on page 2-11.

Logical Disk Layout

Each server has 12 disks. The critical operating system is stored on disks 1 and 2.

[Table 2–7](#) describes how the disks are partitioned.

Table 2–7 Logical Disk Layout

Disk	Description
1 to 2	150 gigabytes (GB) physical and logical partition, mirrored to create two copies, with the Linux operating system, all installed software, NameNode data, and MySQL Database data. The NameNode and MySQL Database data are replicated on two servers for a total of four copies. 2.8 terabytes (TB) HDFS data partition
3 to 12	Single HDFS or Oracle NoSQL Database data partition

Critical and Noncritical Nodes

Critical nodes are required for the cluster to operate normally and provide all services to users. In contrast, the cluster continues to operate with no loss of service when a noncritical node fails.

On single-rack clusters, the critical services are installed initially on the first four nodes of the cluster. The remaining nodes (node05 up to node18) only run noncritical services. If a hardware failure occurs on one of the critical nodes, then the services can be moved to another, noncritical server. For example, if node02 fails, then you might move its critical services node05. [Table 2–1](#) identifies the initial location of services for clusters that are configured on a single rack.

In a multirack cluster, some critical services run on the first server of the second rack. See ["Where Do the Services Run on a Single Rack?"](#) on page 2-11.

Moving a critical node requires that all clients be reconfigured with the address of the new node. The other alternative is to wait for the failed server to be repaired. You must weigh the loss of services against the inconvenience of reconfiguring the clients.

Table 2–8 Critical Service Locations

Node Name	Initial Node Position	Critical Functions
First NameNode	Node01	ZooKeeper ¹ , first NameNode ¹ , NameNode failover controller ¹ , balancer ¹ , puppet master
Second NameNode	Node02	ZooKeeper, second NameNode, NameNode failover controller, MySQL backup server
First ResourceManager Node	Node03	ZooKeeper, first ResourceManager, Cloudera Manager server, MySQL primary server
Second ResourceManager Node	Node04	Second ResourceManager, Oracle Data Integrator agent, Hue, Hive, Oozie

¹ In multirack clusters, this service is initially installed on the first server of the second rack.

First Namenode

One instance of the NameNode initially runs on node01. If this node fails or goes offline (such as a restart), then the second NameNode (node02) automatically takes over to maintain the normal activities of the cluster.

Alternatively, if the second NameNode is already active, it continues without a backup. With only one NameNode, the cluster is vulnerable to failure. The cluster has lost the redundancy needed for automatic failover.

The puppet master also runs on this node. The Mammoth utility uses Puppet, and so you cannot install or reinstall the software if, for example, you must replace a disk drive elsewhere in the rack.

Note: In multirack clusters, the NameNode service is installed on the first server of the *second* rack.

Second NameNode

One instance of the NameNode initially runs on node02. If this node fails, then the function of the NameNode either fails over to the first NameNode (node01) or continues there without a backup. However, the cluster has lost the redundancy needed for automatic failover if the first NameNode also fails.

The MySQL backup database also runs on this node. MySQL Database continues to run, although there is no backup of the master database.

First ResourceManager

One instance of the ResourceManager initially runs on node03. If this node fails or goes offline (such as a restart, then the second ResourceManager (node04) automatically takes over to distribute MapReduce tasks to specific nodes across the cluster.

Alternatively, if the second ResourceManager is already active, it continues without a backup. With only one ResourceManager, the cluster is vulnerable to failure. The cluster has lost the redundancy needed for automatic failover.

These services are also disrupted:

- **Cloudera Manager:** This tool provides central management for the entire CDH cluster. Without this tool, you can still monitor activities using the utilities described in ["Using Hadoop Monitoring Utilities"](#) on page 2-5.
- **MySQL Master Database:** Cloudera Manager, Oracle Data Integrator, Hive, and Oozie use MySQL Database. The data is replicated automatically, but you cannot access it when the master database server is down.

Second ResourceManager

One instance of the ResourceManager initially runs on node04. If this node fails, then the function of the ResourceManager either fails over to the first ResourceManager (node03) or continues there without a backup. However, the cluster has lost the redundancy needed for automatic failover if the first ResourceManager also fails.

These services are also disrupted:

- **Oracle Data Integrator:** This service supports Oracle Data Integrator Application Adapter for Hadoop. You cannot use this connector when the ResourceManager node is down.
- **Hive:** Hive provides a SQL-like interface to data that is stored in HDFS. Most of the Oracle Big Data Connectors can access Hive tables, which are not available if this node fails.
- **Hue:** This administrative tool is not available when the ResourceManager node is down.

- **Oozie:** This workflow and coordination service runs on the ResourceManager node, and is unavailable when the node is down.

Noncritical Nodes

The noncritical nodes (node05 to node18) are optional in that Oracle Big Data Appliance continues to operate with no loss of service if a failure occurs. The NameNode automatically replicates the lost data to always maintain three copies. MapReduce jobs execute on copies of the data stored elsewhere in the cluster. The only loss is in computational power, because there are fewer servers on which to distribute the work.

Stopping and Starting Oracle Big Data Appliance

This section describes how to shut down Oracle Big Data Appliance gracefully and restart it.

- [Prerequisites](#)
- [Stopping Oracle Big Data Appliance](#)
- [Starting Oracle Big Data Appliance](#)

Prerequisites

You must have root access. Passwordless SSH must be set up on the cluster, so that you can use the `dcli` utility.

To ensure that passwordless-ssh is set up:

1. Log in to the first node of the cluster as root.
2. Use a `dcli` command to verify it is working. This command should return the IP address and host name of every node in the cluster:

```
# dcli -C hostname
192.0.2.1: bda1node01.example.com
192.0.2.2: bda1node02.example.com
.
.
.
```

3. If you do not get these results, then set up `dcli` on the cluster:

```
# setup-root-ssh -C
```

See Also: *Oracle Big Data Appliance Owner's Guide* for details about these commands.

Stopping Oracle Big Data Appliance

Follow these procedures to shut down all Oracle Big Data Appliance software and hardware components.

Note: The following services stop automatically when the system shuts down. You do not need to take any action:

- Oracle Enterprise Manager agent
 - Auto Service Request agents
-
-

Task 1 Stopping All Managed Services

Use Cloudera Manager to stop the services it manages, including flume, hbase, hdfs, hive, hue, mapreduce, oozie, and zookeeper.

1. Log in to Cloudera Manager as the admin user.

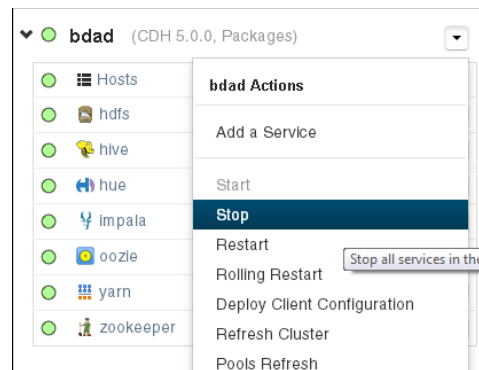
See ["Managing CDH Operations Using Cloudera Manager"](#) on page 2-3.

2. In the Status pane of the opening page, expand the menu for the cluster and click **Stop**, and then click **Stop** again when prompted to confirm. See [Figure 2-9](#).

To navigate to this page, click the **Home** tab, and then the **Status** subtab.

3. On the Command Details page, click **Close** when all processes are stopped.
4. In the same pane under Cloudera Management Services, expand the menu for the mgmt service and click **Stop**.
5. Log out of Cloudera Manager.

Figure 2-9 Stopping HDFS Services



Task 2 Stopping Cloudera Manager Server

Follow this procedure to stop Cloudera Manager Server.

1. Log in as root to the node where Cloudera Manager runs (initially node03).

Note: The remaining tasks presume that you are logged in to a server as root. You can enter the commands from any server by using the `dcli` command. This example runs the `pwd` command on node03 from any node in the cluster:

```
# dcli -c node03 pwd
```

2. Stop the Cloudera Manager server:

```
# service cloudera-scm-server stop
Stopping cloudera-scm-server: [ OK ]
```

3. Verify that the server is stopped:

```
# service cloudera-scm-server status
cloudera-scm-server is stopped
```

After stopping Cloudera Manager, you cannot access it using the web console.

Task 3 Stopping Oracle Data Integrator Agent

If Oracle Data Integrator Application Adapter for Hadoop is installed on the cluster, then stop the agent.

1. Check the status of the Oracle Data Integrator service:

```
# dcli -C service odi-agent status
```

2. Stop the Oracle Data Integrator agent, if it is running:

```
# dcli -C service odi-agent stop
```

3. Ensure that the Oracle Data Integrator service stopped running:

```
# dcli -C service odi-agent status
```

Task 4 Dismounting NFS Directories

All nodes share an NFS directory on node03, and additional directories may also exist. If a server with the NFS directory (/opt/exportdir) is unavailable, then the other servers hang when attempting to shut down. Thus, you must dismount the NFS directories first.

1. Locate any mounted NFS directories:

```
# dcli -C mount | grep sharedir
192.0.2.1: bdalnode03.example.com:/opt/exportdir on /opt/sharedir type nfs
(rw,tcp,soft,intr,timeo=10,retrans=10,addr=192.0.2.3)
192.0.2.2: bdalnode03.example.com:/opt/exportdir on /opt/sharedir type nfs
(rw,tcp,soft,intr,timeo=10,retrans=10,addr=192.0.2.3)
192.0.2.3: /opt/exportdir on /opt/sharedir type none (rw,bind)
.
.
.
```

The sample output shows a shared directory on node03 (192.0.2.3).

2. Dismount the shared directory:

```
# dcli -C umount /opt/sharedir
```

3. Dismount any custom NFS directories.

Task 5 Stopping the Servers

The Linux shutdown -h command powers down individual servers. You can use the dcli -g command to stop multiple servers.

1. Create a file that lists the names or IP addresses of the other servers in the cluster, that is, not including the one you are logged in to.
2. Stop the other servers:

```
# dcli -g filename shutdown -h now
```

For *filename*, enter the name of the file that you created in step 1.

3. Stop the server you are logged in to:

```
# shutdown -h now
```

Task 6 Stopping the InfiniBand and Cisco Switches

To stop the network switches, turn off a PDU or a breaker in the data center. The switches only turn off when power is removed.

The network switches do not have power buttons. They shut down only when power is removed

To stop the switches, turn off all breakers in the two PDUs.

Starting Oracle Big Data Appliance

Follow these procedures to power up the hardware and start all services on Oracle Big Data Appliance.

Task 1 Powering Up Oracle Big Data Appliance

1. Switch on all 12 breakers on both PDUs.
2. Allow 4 to 5 minutes for Oracle ILOM and the Linux operating system to start on the servers.
3. If password-based, on-disk encryption is enabled, then log in and mount the Hadoop directories on those servers:

```
$ mount-hadoop-dirs
Enter password to mount Hadoop directories: password
```

If the servers do not start automatically, then you can start them locally by pressing the power button on the front of the servers, or remotely by using Oracle ILOM. Oracle ILOM has several interfaces, including a command-line interface (CLI) and a web console. Use whichever interface you prefer.

For example, you can log in to the web interface as `root` and start the server from the Remote Power Control page. The URL for Oracle ILOM is the same as for the host, except that it typically has a `-c` or `-ilom` extension. This URL connects to Oracle ILOM for `bda1node4`:

```
http://bda1node04-ilom.example.com
```

Task 2 Starting the HDFS Software Services

Use Cloudera Manager to start all the HDFS services that it controls.

1. Log in as `root` to the node where Cloudera Manager runs (initially `node03`).

Note: The remaining tasks presume that you are logged in to a server as `root`. You can enter the commands from any server by using the `dcli` command. This example runs the `pwd` command on `node03` from any node in the cluster:

```
# dcli -c node03 pwd
```

2. Verify that the Cloudera Manager started automatically on `node03`:

```
# service cloudera-scm-server status
cloudera-scm-server (pid 11399) is running...
```

3. If it is not running, then start it:

```
# service cloudera-scm-server start
```

4. Log in to Cloudera Manager as the `admin` user.

See ["Managing CDH Operations Using Cloudera Manager"](#) on page 2-3.

5. In the Status pane of the opening page, expand the menu for the cluster and click **Start**, and then click **Start** again when prompted to confirm. See [Figure 2–9](#).
To navigate to this page, click the **Home** tab, and then the **Status** subtab.
6. On the Command Details page, click **Close** when all processes are started.
7. In the same pane under Cloudera Management Services, expand the menu for the mgmt service and click **Start**.
8. Log out of Cloudera Manager (optional).

Task 3 Starting Oracle Data Integrator Agent

If Oracle Data Integrator Application Adapter for Hadoop is used on this cluster, then start the agent.

1. Check the status of the agent:

```
# /opt/oracle/odiagent/agent_standalone/oracledi/agent/bin/startcmd.sh  
OdiPingAgent [-AGENT_NAME=agent_name]
```

2. Start the agent:

```
# /opt/oracle/odiagent/agent_standalone/oracledi/agent/bin/agent.sh  
[-NAME=agent_name] [-PORT=port_number]
```

Switching From YARN to MapReduce 1

Oracle Big Data Appliance uses the Yet Another Resource Negotiator (YARN) implementation of MapReduce by default. You have the option of using classic MapReduce (MRv1) instead. You cannot use both implementations in the same cluster; you can activate either the MapReduce or the YARN service.

To switch a cluster to MRv1:

1. Log in to Cloudera Manager as the admin user.
2. Stop the YARN service:
 - a. Locate YARN in the list of services on the Status tab of the Home page.
 - b. Expand the YARN menu and click **Stop**.
3. On the cluster menu, click **Add a Service** to start the Add Service wizard:
 - a. Select MapReduce for the type of service you want to add.
 - b. Select hdfs/zookeeper as a dependency (default).
 - c. Customize the role assignments:

JobTracker: Click the field to display a list of nodes in the cluster. Select the third node.

TaskTracker: For a six-node cluster, keep the TaskTrackers on all nodes (default). For larger clusters, remove the TaskTrackers from the first two nodes.

- d. On the Review Changes page, change the parameter values:

TaskTracker Local Data Directory List: Change the default group and group 1 to `"/u12/hadoop/mapred".." /u01/hadoop/mapred"`.

JobTracker Local Data Directory List: Change the default group to `"/u12/hadoop/mapred".." /u01/hadoop/mapred"`.

- e. Complete the steps of the wizard with no further changes. Click **Finish** to save the configuration and exit.
4. Update the Hive service configuration:
 - a. On the Status tab of the Home page, click **hive** to display the hive page.
 - b. Expand the Configuration submenu and click **View and Edit**.
 - c. Select **mapreduce** as the value of the MapReduce Service property.
 - d. Click **Save Changes**.
5. Repeat step 4 to update the Oozie service configuration to use **mapreduce**.
6. On the Status tab of the Home page, expand the **hive** and **oozie** menus and choose **Restart**.
7. Optional: Expand the **yarn** service menu and choose **Delete**.
 If you retain the **yarn** service, then after every cluster restart, you will see "Memory overcommit validation" warnings, and you must manually stop the **yarn** service.
8. Update the MapReduce service configuration:
 - a. On the Status tab of the Home page, click **mapreduce** to display the **mapreduce** page.
 - b. Expand the Configuration submenu and click **View and Edit**.
 - c. Under Category, expand **TaskTracker Default Group**, and then click **Resource Management**.
 - d. Set the following properties:

Java Heap Size of TaskTracker in Bytes: Reset to the default value of 1 GiB.

Maximum Number of Simultaneous Map Tasks: Set to either 15 for Sun Fire X4270 M2 racks or 20 for all other racks.

Maximum Number of Simultaneous Reduce Tasks: Set to either 10 for Sun Fire X4270 M2 racks or 13 for all other racks.
 - e. Click **Save Changes**.
9. Add overrides for nodes 3 and 4 (or nodes 1 and 2 in a six-node cluster). See "[Map and Reduce Resource Configuration](#)" on page 2-14.
10. Click the **mapreduce1** service to display the **mapreduce** page:
11. Expand the Actions menu and select **Enable High Availability** to start the Enable JobTracker High Availability wizard:
 - a. On the Assign Roles page, select the fourth node (node04) for the Standby JobTracker.
 - b. Complete the steps of the wizard with no further changes. Click **Finish** to save the configuration and exit.
12. Verify that all services in the cluster are healthy with no configuration issues.
13. Reconfigure Perfect Balance for the MRv1 cluster:
 - a. Log in as **root** to a node of the cluster.
 - b. Configure Perfect Balance on all nodes of the cluster:

```
$ dcli -C /opt/oracle/orabalancer-[version]/bin/configure.sh
```

See *Oracle Big Data Appliance Owner's Guide* for the `dcli` command.

Security on Oracle Big Data Appliance

You can take precautions to prevent unauthorized use of the software and data on Oracle Big Data Appliance.

This section contains these topics:

- [About Predefined Users and Groups](#)
- [About User Authentication](#)
- [About Fine-Grained Authorization](#)
- [About On-Disk Encryption](#)
- [Port Numbers Used on Oracle Big Data Appliance](#)
- [About Puppet Security](#)

About Predefined Users and Groups

Every open-source package installed on Oracle Big Data Appliance creates one or more users and groups. Most of these users do not have login privileges, shells, or home directories. They are used by daemons and are not intended as an interface for individual users. For example, Hadoop operates as the `hdfs` user, MapReduce operates as `mapred`, and Hive operates as `hive`.

You can use the `oracle` identity to run Hadoop and Hive jobs immediately after the Oracle Big Data Appliance software is installed. This user account has login privileges, a shell, and a home directory.

Oracle NoSQL Database and Oracle Data Integrator run as the `oracle` user. Its primary group is `oinstall`.

Note: Do not delete, re-create, or modify the users that are created during installation, because they are required for the software to operate.

Table 2–9 identifies the operating system users and groups that are created automatically during installation of Oracle Big Data Appliance software for use by CDH components and other software packages.

Table 2–9 Operating System Users and Groups

User Name	Group	Used By	Login Rights
flume	flume	Apache Flume parent and nodes	No
hbase	hbase	Apache HBase processes	No
hdfs	hadoop	NameNode , DataNode	No
hive	hive	Hive metastore and server processes	No
hue	hue	Hue processes	No
mapred	hadoop	ResourceManager, NodeManager, Hive Thrift daemon	Yes
mysql	mysql	MySQL server	Yes

Table 2–9 (Cont.) Operating System Users and Groups

User Name	Group	Used By	Login Rights
oozie	oozie	Oozie server	No
oracle	dba, oinstall	Oracle NoSQL Database, Oracle Loader for Hadoop, Oracle Data Integrator, and the Oracle DBA	Yes
puppet	puppet	Puppet parent (puppet nodes run as root)	No
sqoop	sqoop	Apache Sqoop metastore	No
svctag		Auto Service Request	No
zookeeper	zookeeper	ZooKeeper processes	No

About User Authentication

Oracle Big Data Appliance supports Kerberos security as a software installation option. See [Chapter 3](#) for details about setting up clients and users to access a Kerberos-protected cluster.

About Fine-Grained Authorization

The typical authorization model on Hadoop is at the HDFS file level, such that users either have access to all of the data in the file or none. In contrast, Apache Sentry integrates with the Hive and Impala SQL-query engines to provide fine-grained authorization to data and metadata stored in Hadoop.

Oracle Big Data Appliance automatically configures Sentry during software installation, beginning with Mammoth utility version 2.5.

See Also:

- Cloudera Manager Help
- *Managing Clusters with Cloudera Manager* at <https://www.cloudera.com/content/cloudera-content/cloudera-docs/CM4Ent/latest/Cloudera-Manager-Managing-Clusters/Managing-Clusters-with-Cloudera-Manager.html>

About On-Disk Encryption

On-disk encryption protects data that is at rest on disk. When on-disk encryption is enabled, Oracle Big Data Appliance automatically encrypts and decrypts data stored on disk. On-disk encryption does not affect user access to Hadoop data, although it can have a minor impact on performance.

Password-based encryption encodes Hadoop data based on a password, which is the same for all servers in a cluster. You can change the password at any time by using the `mammoth-reconfig update` command. See *Oracle Big Data Appliance Owner's Guide*.

If a disk is removed from a server, then the encrypted data remains protected until you install the disk in a server (the same server or a different one), startup the server, and provide the password. If a server is powered off and removed from an Oracle Big Data Appliance rack, then the encrypted data remains protected until you restart server and provide the password. You must enter the password after every startup of every server to enable access to the data. See ["Starting Oracle Big Data Appliance"](#) on page 2-21.

On-disk encryption is an option that you can select during the initial installation of the software by the Mammoth utility. You can also enable or disable on-disk encryption at any time by using either the `mammoth-reconfig` or `bdaccli` utilities.

See Also: *Oracle Big Data Appliance Owner's Guide*

Port Numbers Used on Oracle Big Data Appliance

Table 2–10 identifies the port numbers that might be used in addition to those used by CDH.

To view the ports used on a particular server:

1. In Cloudera Manager, click the **Hosts** tab at the top of the page to display the Hosts page.
2. In the Name column, click a server link to see its detail page.
3. Scroll down to the Ports section.

See Also: For the full list of CDH port numbers, go to the Cloudera website at

http://www.cloudera.com/content/cloudera-content/cloudera-docs/CM5/latest/Cloudera-Manager-Installation-Guide/cm5ig_config_ports.html

Table 2–10 Oracle Big Data Appliance Port Numbers

Service	Port
Automated Service Monitor (ASM)	30920
HBase master service (node01)	60010
MySQL Database	3306
Oracle Data Integrator Agent	20910
Oracle NoSQL Database administration	5001
Oracle NoSQL Database processes	5010 to 5020
Oracle NoSQL Database registration	5000
Port map	111
Puppet master service	8140
Puppet node service	8139
rpc.statd	668
ssh	22
xinetd (service tag)	6481

About Puppet Security

The puppet node service (`puppetd`) runs continuously as `root` on all servers. It listens on port 8139 for "kick" requests, which trigger it to request updates from the puppet master. It does not receive updates on this port.

The puppet master service (`puppetmasterd`) runs continuously as the puppet user on the first server of the primary Oracle Big Data Appliance rack. It listens on port 8140 for requests to push updates to puppet nodes.

The puppet nodes generate and send certificates to the puppet master to register initially during installation of the software. For updates to the software, the puppet master signals ("kicks") the puppet nodes, which then request all configuration changes from the puppet master node that they are registered with.

The puppet master sends updates only to puppet nodes that have known, valid certificates. Puppet nodes only accept updates from the puppet master host name they initially registered with. Because Oracle Big Data Appliance uses an internal network for communication within the rack, the puppet master host name resolves using `/etc/hosts` to an internal, private IP address.

Auditing Oracle Big Data Appliance

You can use Oracle Audit Vault and Database Firewall to create and monitor the audit trails for HDFS and MapReduce on Oracle Big Data Appliance.

This section describes the Oracle Big Data Appliance plug-in:

- [About Oracle Audit Vault and Database Firewall](#)
- [Setting Up the Oracle Big Data Appliance Plug-in](#)
- [Monitoring Oracle Big Data Appliance](#)

About Oracle Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall secures databases and other critical components of IT infrastructure in these key ways:

- Provides an integrated auditing platform for your enterprise.
- Captures activity on Oracle Database, Oracle Big Data Appliance, operating systems, directories, file systems, and so forth.
- Makes the auditing information available in a single reporting framework so that you can understand the activities across the enterprise. You do not need to monitor each system individually; you can view your computer infrastructure as a whole.

Audit Vault Server provides a web-based, graphic user interface for both administrators and auditors.

You can configure CDH/Hadoop clusters on Oracle Big Data Appliance as secured targets. The Audit Vault plug-in on Oracle Big Data Appliance collects audit and logging data from these services:

- **HDFS:** Who makes changes to the file system.
- **Hive DDL:** Who makes Hive database changes.
- **MapReduce:** Who runs MapReduce jobs that correspond to file access.
- **Oozie workflows:** Who runs workflow activities.

The Audit Vault plug-in is an installation option. The Mammoth utility automatically configures monitoring on Oracle Big Data Appliance as part of the software installation process.

See Also: For more information about Oracle Audit Vault and Database Firewall:

<http://www.oracle.com/technetwork/database/database-technologies/audit-vault-and-database-firewall/overview/index.html>

Setting Up the Oracle Big Data Appliance Plug-in

The Mammoth utility on Oracle Big Data Appliance performs all the steps needed to setup the plug-in, using information that you provide.

To set up the Audit Vault plug-in for Oracle Big Data Appliance:

1. Ensure that Oracle Audit Vault and Database Firewall Server Release 12.1.1 is up and running on the same network as Oracle Big Data Appliance.

See Also: *Oracle Audit Vault and Database Firewall Installation Guide*

2. Complete the Audit Vault Plug-in section of Oracle Big Data Appliance Configuration Generation Utility.
3. Install the Oracle Big Data Appliance software using the Mammoth utility. An Oracle representative typically performs this step.

You can also add the plug-in at a later time using either `bdaccli` or `mammoth-reconfig`. See *Oracle Big Data Appliance Owner's Guide*.

When the software installation is complete, the Audit Vault plug-in is installed on Oracle Big Data Appliance, and Oracle Audit Vault and Database Firewall is collecting its audit information. You do not need to perform any other installation steps.

See Also: *Oracle Big Data Appliance Owner's Guide* for using Oracle Big Data Appliance Configuration Generation Utility

Monitoring Oracle Big Data Appliance

After installing the plug-in, you can monitor Oracle Big Data Appliance the same as any other secured target. Audit Vault Server collects activity reports automatically.

The following procedure describes one type of monitoring activity.

To view an Oracle Big Data Appliance activity report:

1. Log in to Audit Vault Server as an auditor.
2. Click the **Reports** tab.
3. Under Built-in Reports, click **Audit Reports**.
4. To browse all activities, in the Activity Reports list, click the **Browse report data** icon for All Activity.
5. Add or remove the filters to list the events.

Event names include ACCESS, CREATE, DELETE, and OPEN.

6. Click the **Single row view** icon in the first column to see a detailed report.

[Figure 2–10](#) shows the beginning of an activity report, which records access to a Hadoop sequence file.

Figure 2–10 Activity Report in Audit Vault Server

All Activity Report	
Report View	< Row 0 of <input type="checkbox"/> Exclude Null Values <input type="checkbox"/> Displayed Columns
<div>Secured Target</div> <div> Secured Target Name bda1node02 Secured Target Type Oracle BDA Hadoop Cluster Service Name Policy Name Class Other </div>	
<div>Event</div> <div> Server Time 8/20/2013 3:09:51 PM Event Time 8/14/2013 9:28:06 AM User Name mapred(AUTH:SIMPLE) Event Status SUCCESS Error Code Error Message Event Name ACCESS Command Class SELECT Action Taken Threat Severity Log Cause Location Audit File </div>	
<div>Target</div> <div> Target Type Target Object /tmp/mapred/system/seq-000000000003 Target Owner </div>	

See Also: *Oracle Audit Vault and Database Firewall Auditor's Guide*

Collecting Diagnostic Information for Oracle Customer Support

If you need help from Oracle Support to troubleshoot CDH issues, then you should first collect diagnostic information using the `bdadiag` utility with the `cm` option.

To collect diagnostic information:

1. Log in to an Oracle Big Data Appliance server as `root`.
2. Run `bdadiag` with at least the `cm` option. You can include additional options on the command line as appropriate. See the *Oracle Big Data Appliance Owner's Guide* for a complete description of the `bdadiag` syntax.

```
# bdadiag cm
```

The command output identifies the name and the location of the diagnostic file.

3. Go to My Oracle Support at <http://support.oracle.com>.
4. Open a Service Request (SR) if you have not already done so.
5. Upload the `bz2` file into the SR. If the file is too large, then upload it to `sftp.oracle.com`, as described in the next procedure.

To upload the diagnostics to `ftp.oracle.com`:

1. Open an SFTP client and connect to `sftp.oracle.com`. Specify port 2021 and remote directory `/support/incoming/target`, where *target* is the folder name given to you by Oracle Support.

See [Example 2–1](#) if you are using a command-line SFTP client.

2. Log in with your Oracle Single Sign-on account and password.
3. Upload the diagnostic file to the new directory.
4. Update the SR with the full path and the file name.

[Example 2–1](#) shows the commands to upload the diagnostics using the SFTP command interface.

Example 2–1 Uploading Diagnostics Using FTP

```
$ sftp -o port=2021 my.user.name@oracle.com@sftp.oracle.com
Connecting to sftp.oracle.com...
.
.
.
Enter password for my.user.name@oracle.com
Password: password
sftp> cd support/incoming/SR123456
sftp> put /tmp/bdadiag_bda1node01_1216FM5497_2013_07_18_07_33.tar.bz2
Uploading bdadiag_bda1node01_1216FM5497_2013_07_18_07_33.tar.bz2 to
//support/incoming/create_table.sql
bdadiag_bda1node01_1216FM5497_2013_07_18_07_33.tar.bz2 to support/incoming/create_
table.sql          100% 311      0.3KB/s   00:00
sftp> exit
$
```

See Also: My Oracle Support Note 549180.1 at

<http://support.oracle.com>

Supporting User Access to Oracle Big Data Appliance

This chapter describes how you can support users who run MapReduce jobs on Oracle Big Data Appliance or using Oracle Big Data Connectors. It contains these sections:

- [About Accessing a Kerberos-Secured Cluster](#)
- [Providing Remote Client Access to CDH](#)
- [Providing Remote Client Access to Hive](#)
- [Managing User Accounts](#)
- [Recovering Deleted Files](#)

About Accessing a Kerberos-Secured Cluster

Apache Hadoop is not an inherently secure system. It is protected only by network security. After a connection is established, a client has full access to the system.

To counterbalance this open environment, Oracle Big Data Appliance supports Kerberos security as a software installation option. Kerberos is a network authentication protocol that helps prevent malicious impersonation.

CDH provides these securities when configured to use Kerberos:

- The CDH master nodes, NameNodes, and JournalNodes resolve the group name so that users cannot manipulate their group memberships.
- Map tasks run under the identity of the user who submitted the job.
- Authorization mechanisms in HDFS and MapReduce help control user access to data.

If the Oracle Big Data Appliance cluster is secured with Kerberos, then you must take additional steps to authenticate a CDH client and individual users, as described in this chapter. Users must know their Kerberos user name, password, and realm.

[Table 3–1](#) describes some frequently used Kerberos commands. For more information, see the MIT Kerberos documentation.

Table 3–1 Kerberos User Commands

Command	Description
<code>kinit userid@realm</code>	Obtains a Kerberos ticket.
<code>klist</code>	Lists a Kerberos ticket if you have one already.
<code>kdestroy</code>	Invalidates a ticket before it expires.

Table 3–1 (Cont.) Kerberos User Commands

Command	Description
<code>kpasswd userid@realm</code>	Changes your password.

See Also:

- MIT Kerberos Documentation at <http://web.mit.edu/kerberos/krb5-latest/doc/>
- *CDH4 Security Guide* at <http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH5/latest/CDH5-Security-Guide/CDH5-Security-Guide.html>

Providing Remote Client Access to CDH

Oracle Big Data Appliance supports full local access to all commands and utilities in Cloudera's Distribution including Apache Hadoop (CDH).

You can use a browser on any computer that has access to the client network of Oracle Big Data Appliance to access Cloudera Manager, Hadoop Map/Reduce Administration, the Hadoop Task Tracker interface, and other browser-based Hadoop tools.

To issue Hadoop commands remotely, however, you must connect from a system configured as a CDH client with access to the Oracle Big Data Appliance client network. This section explains how to set up a computer so that you can access HDFS and submit MapReduce jobs on Oracle Big Data Appliance.

See Also: My Oracle Support ID 1506203.1

Prerequisites

Ensure that you have met the following prerequisites:

- You must have these access privileges:
 - Sudo access to the client system
 - Login access to Cloudera Manager

If you do not have these privileges, then contact your system administrator for help.
- The client system must run an operating system that Cloudera supports for CDH5. See the *Cloudera CDH5 Installation Guide* at

http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH5/latest/CDH5-Requirements-and-Supported-Versions/cdhrsv_os.html

- The client system must run Oracle JDK 1.7.0_25 or later.

To verify the version, use this command:

```
$ java -version
java version "1.7.0_25"
Java(TM) SE Runtime Environment (build 1.7.0_25-b15)
Java HotSpot(TM) 64-Bit Server VM (build 23.25-b01, mixed mode)
```

Installing CDH on Oracle Exadata Database Machine

When you use Oracle Exadata Database Machine as the client, you can use the RPM files on Oracle Big Data Appliance when both engineered systems use the same operating system (Oracle Linux 5.x or 6.x). Copying the files across the local network is faster than downloading them from the Cloudera website.

Note: In the following steps, replace *version_number* with the missing portion of the file name, such as `cdh5.0.0+28-1.cdh5.0.0.p0.36.el6`.

To install a CDH client on Oracle Exadata Database Machine:

1. Log into an Exadata database server.
2. Verify that Hadoop is not installed on your Exadata system:

```
rpm -qa | grep hadoop
```

3. If the `rpm` command returns a value, then remove the existing Hadoop software:

```
rpm -e hadoop_rpm
```

4. Copy the following Linux RPMs to the database server from the first server of Oracle Big Data Appliance. The RPMs are located in the `/opt/oracle/BDAMammoth/bdarepo/RPMS/x86_64` directory.

- `ed-version_number.x86_64.rpm`
- `m4-version_number.x86_64.rpm`
- `nc-version_number.x86_64.rpm`
- `redhat-lsb-version_number.x86_64.rpm`

5. Install the Oracle Linux RPMs from Step 4 on all database nodes. For example:

```
sudo yum --nogpgcheck localinstall ed-1.1-3.3.el6.x86_64.rpm
sudo yum --nogpgcheck localinstall m4-1.4.13-5.el6.x86_64.rpm
sudo yum --nogpgcheck localinstall nc-1.84-22.el6.x86_64.rpm
sudo yum --nogpgcheck localinstall redhat-lsb-4.0-7.0.1.el6.x86_64.rpm
```

Be sure to install the Oracle Linux RPMs before installing the CDH RPMs.

6. Copy the following CDH RPM from the `/opt/oracle/BDAMammoth/bdarepo/RPMS/noarch` directory:

- `bigtop-utils-version_number.noarch.rpm`

7. Copy the following CDH RPMs from the `/opt/oracle/BDAMammoth/bdarepo/RPMS/x86_64` directory:

- `zookeeper-version_number.noarch.rpm`
- `hadoop-version_number.x86_64.rpm`
- `bigtop-jsvc-version_number.x86_64.rpm`
- `hadoop-hdfs-version_number.x86_64.rpm`
- `hadoop-0.20-mapreduce-version_number.x86_64.rpm`
- `hadoop-yarn-version_number.x86_64.rpm`
- `hadoop-mapreduce-version_number.x86_64.rpm`
- `hadoop-client-version_number.x86_64.rpm`

8. Install the CDH RPMs in the exact order shown in Steps 6 and 7 on all database servers. For example:

```
rpm -ihv bigtop-utils-0.7.0+cdh5.0.0+0-1.cdh5.0.0.p0.36.el6.noarch.rpm
rpm -ihv zookeeper-3.4.5+cdh5.0.0+28-1.cdh5.0.0.p0.36.el6.x86_64.rpm
rpm -ihv hadoop-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
rpm -ihv bigtop-jsvc-0.6.0+cdh5.0.0+427-1.cdh5.0.0.p0.34.el6.x86_64.rpm
rpm -ihv hadoop-hdfs-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
rpm -ihv hadoop-0.20-mapreduce-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
rpm -ihv hadoop-yarn-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
rpm -ihv hadoop-mapreduce-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
rpm -ihv hadoop-client-2.3.0+cdh5.0.0+548-1.cdh5.0.0.p0.69.el6.x86_64.rpm
```

9. Configure the CDH client. See ["Configuring a CDH Client for an Unsecured Cluster"](#) on page 3-4.

Installing a CDH Client on Any Supported Operating System

To install a CDH client on any operating system identified as supported by Cloudera, follow these instructions.

To install the CDH client software:

1. Log in to the client system.
2. If an earlier version of Hadoop is already installed, then remove it.

See the Cloudera documentation for removing CDH4 at

http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH5/latest/CDH5-Installation-Guide/cdh5ig_to_cdh5_upgrade.html?scroll=topic_6_3_4_unique_1

3. Download the CDH5 software from the Cloudera website:

```
$ wget
http://archive-primary.cloudera.com/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.0.tar.gz
```

4. Decompress the file into a permanent location, which will be the Hadoop home directory. The following command unzips the files into `hadoop-2.3.0-cdh5.0.0` in the current directory:

```
tar -xvzf hadoop-2.3.0-cdh5.0.0.tar.gz
```

5. Configure the CDH client. See ["Configuring a CDH Client for an Unsecured Cluster"](#) on page 3-4.

Configuring a CDH Client for an Unsecured Cluster

After installing CDH, you must configure it for use with Oracle Big Data Appliance.

To configure the Hadoop client:

1. Log in to the client system and download the MapReduce client configuration from Cloudera Manager. In this example, Cloudera Manager listens on port 7180 (the default) of `bda01node03.example.com`, and the configuration is stored in a file named `yarn-conf.zip`.

```
$ wget -O yarn-conf.zip
http://bda01node03.example.com:7180/cmfservices/3/client-config
```

2. Unzip `mapreduce-config.zip` into a permanent location on the client system.

```
$ unzip yarn-config.zip
Archive:  yarn-config.zip
  inflating: yarn-conf/hadoop-env.sh
  inflating: yarn-conf/hdfs-site.xml
  inflating: yarn-conf/core-site.xml
  inflating: yarn-conf/mapred-site.xml
  inflating: yarn-conf/log4j.properties
  inflating: yarn-conf/yarn-site.xml
```

All files are stored in a subdirectory named `yarn-config`.

3. Open `hadoop-env.sh` in a text editor and set the environment variables to the actual paths on your system:

```
export HADOOP_HOME=hadoop-home-dir/share/hadoop/mapreduce1
export HADOOP_MAPRED_HOME=$HADOOP_HOME
export HADOOP_CONF_DIR=yarn-conf-dir
export JAVA_HOME=/usr/java/version
ln -s $HADOOP_HOME/../../../../bin-mapreduce1 $HADOOP_HOME/bin
ln -s $HADOOP_HOME/../../../../etc/hadoop-mapreduce1 $HADOOP_HOME/conf
rm -f $HADOOP_HOME/lib/slf4j-log4j.jar
alias hadoop=$HADOOP_HOME/bin/hadoop
alias hdfs=$HADOOP_HOME/../../../../bin/hdfs
```

4. Make a backup copy of the Hadoop configuration files:

```
# cp /full_path/yarn-conf /full_path/yarn-conf-bak
```

5. Overwrite the existing configuration files with the downloaded configuration files in Step 2.

```
# cd /full_path/yarn-conf
# cp * /usr/lib/hadoop/conf
```

Configuring a CDH Client for a Kerberos-Secured Cluster

Follow these steps to enable the CDH client to work with a secure CDH cluster.

To configure a CDH client for Kerberos:

1. Log in to the system where you created the CDH client.
2. Install the Java Cryptography Extension Unlimited Strength Jurisdiction Policy Files:

- a. Download the files for your Java version:

Java 6:

<http://www.oracle.com/technetwork/java/javase/downloads/jce-6-download-429243.html>

Java 7:

<http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>

- b. Decompress the downloaded file. This example unzips JCE-7:

```
$ unzip UnlimitedJCEPolicyJDK7.zip
Archive:  UnlimitedJCEPolicyJDK7.zip
  creating: UnlimitedJCEPolicy/
  inflating: UnlimitedJCEPolicy/US_export_policy.jar
  inflating: UnlimitedJCEPolicy/local_policy.jar
  inflating: UnlimitedJCEPolicy/README.txt
```

Note: The JCE-6 files unzip into a directory named `jce` instead of `UnlimitedJCEPolicy`.

- c. Copy the unzipped files into the Java security directory. For example:

```
$ cp UnlimitedJCEPolicy/* /usr/java/latest/jre/lib/security/
```

3. Follow the steps for configuring an unsecured client.

See ["Configuring a CDH Client for an Unsecured Cluster"](#) on page 3-4.

4. Ensure that you have a user ID on the CDH cluster that had been added to the Kerberos realm.

See ["Creating Hadoop Cluster Users"](#) on page 3-9.

5. On the CDH client system, create a file named `krb5.conf` in the `$HADOOP_CONF_DIR` directory. Enter configuration settings like the following, using values appropriate for your installation for the server names, domain, and realm:

```
[libdefaults]
    default_realm = EXAMPLE.COM
    dns_lookup_realm = false
    dns_lookup_kdc = false
    clockskew = 3600
    ticket_lifetime = 24h
    renew_lifetime = 7d
    forwardable = true

[realms]
    EXAMPLE.COM = {
        kdc = bda01node01.example:88
        admin_server = bda01node07:749
        default_domain = example.com
    }

[domain_realm]
    .com = EXAMPLE.COM
```

6. Activate the new configuration file:

```
export KRB5_CONFIG=$HADOOP_CONF_DIR/krb5.conf
export HADOOP_OPTS="-Djava.security.krb5.conf=$HADOOP_CONF_DIR/krb5.conf "
export KRB5CCNAME=$HADOOP_CONF_DIR/krb5cc_$USER
```

7. Verify that you have access to the Oracle Big Data Appliance cluster.

See ["Verifying Access to a Cluster from the CDH Client"](#) on page 3-6.

Verifying Access to a Cluster from the CDH Client

Follow this procedure to ensure that you have access to the Oracle Big Data Appliance cluster.

To verify cluster access:

1. To access a Kerberos-protected CDH cluster, first obtain a ticket granting ticket (TGT):

```
$ kinit userid@realm
```

2. Verify that you can access HDFS on Oracle Big Data Appliance from the client, by entering a simple Hadoop file system command like the following:

```
$ hadoop fs -ls /user
Found 6 items
```

```
drwxr-xr-x - jdoe      hadoop      0 2014-04-03 00:08 /user/jdoe
drwxrwxrwx - mapred    hadoop      0 2014-04-02 23:25 /user/history
drwxr-xr-x - hive      supergroup  0 2014-04-02 23:27 /user/hive
drwxrwxr-x - impala    impala      0 2014-04-03 10:45 /user/impala
drwxr-xr-x - oozie     hadoop      0 2014-04-02 23:27 /user/oozie
drwxr-xr-x - oracle     hadoop      0 2014-04-03 11:49 /user/oracle
```

Check the output for HDFS users defined on Oracle Big Data Appliance, and not on the client system. You should see the same results as you would after entering the command directly on Oracle Big Data Appliance.

3. Submit a MapReduce job. You must be logged in to the client system under the same user name as your HDFS user name on Oracle Big Data Appliance.

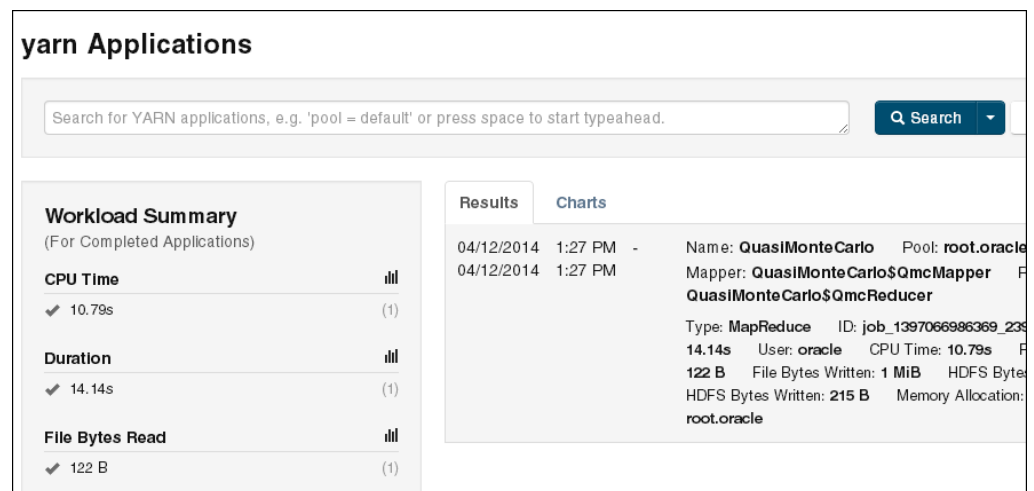
The following example calculates the value of π :

```
$ hadoop jar
/usr/lib/hadoop-mapreduce/hadoop-mapreduce-examples-2.3.0-cdh5.0.0.jar pi 10
1000000
Number of Maps = 10
Samples per Map = 1000000
Wrote input for Map #0
Wrote input for Map #1
.
.
.
Job Finished in 12.403 seconds
Estimated value of Pi is 3.14158440000000000000
```

4. Use Cloudera Manager to verify that the job ran on Oracle Big Data Appliance instead of the local system. Select **mapreduce Jobs** from the Activities menu for a list of jobs.

Figure 3–1 shows the job created by the previous example.

Figure 3–1 Monitoring a YARN Job in Cloudera Manager



Providing Remote Client Access to Hive

Follow this procedure to provide remote client access to Hive.

To set up a Hive client:

1. Set up a CDH client. See ["Providing Remote Client Access to CDH"](#) on page 3-2.
2. Log in to the client system and download the Hive client configuration from Cloudera Manager. In this example, Cloudera Manager listens on port 7180 (the default) of `bda01node03.example.com`, and the configuration is stored in a file named `hive-conf.zip`.

```
$ wget -O hive-conf.zip
http://bda01node03.example.com:7180/cm/services/5/client-config
Length: 1283 (1.3K) [application/zip]
Saving to: 'hive-conf.zip'

100%[=====>] 1,283      --.-K/s   in 0.001s

2014-05-15 08:19:06 (2.17 MB/s) - 'hive-conf.zip' saved [1283/1283]
```

3. Unzip the file into a permanent installation directory, which will be the Hive configuration directory:

```
$ unzip hive-conf.zip
Archive:  hive-conf.zip
  inflating: hive-conf/hive-env.sh
  inflating: hive-conf/hive-site.xml
```

4. Download the Hive software from the Cloudera website:

```
$ wget http://archive.cloudera.com/cdh5/cdh/5/hive-0.12.0-cdh5.0.0.tar.gz
Length: 49637596 (47M) [application/x-gzip]
Saving to: 'hive-0.12.0-cdh5.0.0.tar.gz'

100%[=====>] 49,637,596   839K/s   in 47s

2014-05-15 08:22:18 (1.02 MB/s) - 'hive-0.12.0-cdh5.0.0.tar.gz' saved
[49637596/49637596]
```

5. Decompress the file into a permanent installation directory, which will be the Hive home directory. The following command unzips the files into the current directory in a subdirectory named `hive-0.12.0-cdh5.0.0`:

```
$ tar -xvzf hive-0.12.0-cdh5.0.0.tar.gz
hive-0.12.0-cdh5.0.0/
hive-0.12.0-cdh5.0.0/examples/
.
```

6. Set the following variables, replacing *hive-home-dir* and *hive-conf-dir* with the directories you created in steps 3 and 5.

```
export HIVE_HOME=hive-home-dir
export HIVE_CONF_DIR=hive-conf-dir
alias hive=$HIVE_HOME/bin/hive
```

The following steps test whether you successfully set up a Hive client.

To verify Hive access:

1. To access a Kerberos-protected CDH cluster, first obtain a ticket granting ticket (TGT):


```
$ kinit userid@realm
```

2. Open the Hive console:

```
$ hive
Logging initialized using configuration in
jar:file:/usr/lib/hive/lib/hive-common-0.12.0-cdh5.0.0.jar!/hive-log4j.properties
Hive history file=/tmp/oracle/hive_job_log_
e10527ee-9637-4c08-9559-a2e5cea6cef1_831268640.txt
hive>
```

3. List all tables:

```
hive> show tables;
OK
src
```

Managing User Accounts

This section describes how to create users who can access HDFS, MapReduce, and Hive. It contains the following topics:

- [Creating Hadoop Cluster Users](#)
- [Providing User Login Privileges \(Optional\)](#)

Creating Hadoop Cluster Users

When creating user accounts, define them as follows:

- To run MapReduce jobs, users must either be in the `hadoop` group or be granted the equivalent permissions.
- To create and modify tables in Hive, users must either be in the `hive` group or be granted the equivalent permissions.
- To create Hue users, open Hue in a browser and click the User Admin icon. See ["Using Cloudera Hue to Interact With Hadoop"](#) on page 2-7.

Creating Users on an Unsecured Cluster

To create a user on an unsecured Hadoop cluster:

1. Open an ssh connection as the `root` user to a noncritical node (node04 to node18).
2. Create the user's home directory:

```
# sudo -u hdfs hadoop fs -mkdir /user/user_name
```

You use `sudo` because the HDFS super user is `hdfs` (not `root`).

3. Change the ownership of the directory:

```
# sudo -u hdfs hadoop fs -chown user_name:hadoop /user/user_name
```

4. Verify that the directory is set up correctly:

```
# hadoop fs -ls /user
```

5. Create the operating system user across all nodes in the cluster:

```
# dcli useradd -G hadoop,hive[,group_name...] -m user_name
```

In this syntax, replace *group_name* with an existing group and *user_name* with the new name.

6. Verify that the operating system user belongs to the correct groups:

```
# dcli id user_name
```

7. Verify that the user's home directory was created on all nodes:

```
# dcli ls /home | grep user_name
```

[Example 3-1](#) creates a user named `jdoe` with a primary group of `hadoop` and an addition group of `hive`.

Example 3-1 Creating a Hadoop User

```
# sudo -u hdfs hadoop fs -mkdir /user/jdoe
# sudo -u hdfs hadoop fs -chown jdoe:hadoop /user/jdoe
# hadoop fs -ls /user
Found 5 items
drwx----- - hdfs      supergroup          0 2013-01-16 13:50 /user/hdfs
drwxr-xr-x - hive      supergroup          0 2013-01-16 12:58 /user/hive
drwxr-xr-x - jdoe      jdoe                0 2013-01-18 14:04 /user/jdoe
drwxr-xr-x - oozie     hadoop              0 2013-01-16 13:01 /user/oozie
drwxr-xr-x - oracle    hadoop              0 2013-01-16 13:01 /user/oracle
# dcli useradd -G hadoop,hive -m jdoe
# dcli id jdoe
bda1node01: uid=1001(jdoe) gid=1003(jdoe) groups=1003(jdoe),127(hive),123(hadoop)
bda1node02: uid=1001(jdoe) gid=1003(jdoe) groups=1003(jdoe),123(hadoop),127(hive)
bda1node03: uid=1001(jdoe) gid=1003(jdoe) groups=1003(jdoe),123(hadoop),127(hive)
.
.
.
# dcli ls /home | grep jdoe
bda1node01: jdoe
bda1node02: jdoe
bda1node03: jdoe
```

Creating Users on a Secured Cluster

To create a user on a Kerberos-secured cluster:

1. Connect to Kerberos as the HDFS principal and execute the following commands, replacing `jdoe` with the actual user name:

```
hdfs dfs -mkdir /user/jdoe
hdfs dfs -chown jdoe /user/jdoe
dcli -C useradd -G hadoop,hive -m jdoe
hash=$(echo "hadoop" | openssl passwd -1 -stdin)
dcli -C "usermod --pass='$hash' jdoe"
```

2. Log in to the key distribution center (KDC) and add a principal for the user. In the following example, replace `jdoe`, `bda01node01`, and `example.com` with the correct user name, server name, domain, and realm.

```
ssh -l root bda01node01.example.com kadmin.local
add_principal user_name@EXAMPLE.COM
```

Providing User Login Privileges (Optional)

Users do not need login privileges on Oracle Big Data Appliance to run MapReduce jobs from a remote client. However, for those who want to log in to Oracle Big Data Appliance, you must set a password. You can set or reset a password the same way.

To set a user password across all Oracle Big Data Appliance servers:

1. Create a Hadoop cluster user as described in ["Creating Hadoop Cluster Users"](#) on page 3-9..
2. Confirm that the user does not have a password:

```
# dcli passwd -S user_name
bda1node01.example.com: jdoe NP 2013-01-22 0 99999 7 -1 (Empty password.)
bda1node02.example.com: jdoe NP 2013-01-22 0 99999 7 -1 (Empty password.)
bda1node03.example.com: jdoe NP 2013-01-22 0 99999 7 -1 (Empty password.)
```

If the output shows either "Empty password" or "Password locked," then you must set a password.

3. Set the password:

```
hash=$(echo 'password' | openssl passwd -1 -stdin); dcli "usermod
--pass='$hash' user_name"
```

4. Confirm that the password is set across all servers:

```
# dcli passwd -S user_name
bda1node01.example.com: jdoe PS 2013-01-24 0 99999 7 -1 (Password set, MD5
crypt.)
bda1node02.example.com: jdoe PS 2013-01-24 0 99999 7 -1 (Password set, MD5
crypt.)
bda1node03.example.com: jdoe PS 2013-01-24 0 99999 7 -1 (Password set, MD5
crypt.)
```

See Also:

- *Oracle Big Data Appliance Owner's Guide* for information about dcli.
- The Linux man page for the full syntax of the useradd command.

Recovering Deleted Files

CDH provides an optional trash facility, so that a deleted file or directory is moved to a trash directory for a set period, instead of being deleted immediately from the system. By default, the trash facility is enabled for HDFS and all HDFS clients.

Restoring Files from the Trash

When the trash facility is enabled, you can easily restore files that were previously deleted.

To restore a file from the trash directory:

1. Check that the deleted file is in the trash. The following example checks for files deleted by the oracle user:

```
$ hadoop fs -ls .Trash/Current/user/oracle
Found 1 items
-rw-r--r-- 3 oracle hadoop 242510990 2012-08-31 11:20
```

```
/user/oracle/.Trash/Current/user/oracle/ontime_s.dat
```

2. Move or copy the file to its previous location. The following example moves `ontime_s.dat` from the trash to the HDFS `/user/oracle` directory.

```
$ hadoop fs -mv .Trash/Current/user/oracle/ontime_s.dat /user/oracle/ontime_s.dat
```

Changing the Trash Interval

The **trash interval** is the minimum number of minutes that a file remains in the trash directory before being deleted permanently from the system. The default value is 1 day (24 hours).

To change the trash interval:

1. Open Cloudera Manager. See ["Managing CDH Operations Using Cloudera Manager"](#) on page 2-3.
2. On the Home page under Status, click **hdfs**.
3. On the hdfs page, click the Configuration subtab, and then select **View and Edit**.
4. Search for or scroll down to the Filesystem Trash Interval property under NameNode Default Group. See [Figure 3-2](#).
5. Click the current value, and enter a new value in the pop-up form.
6. Click **Save Changes**.
7. Expand the Actions menu at the top of the page and choose **Restart**.
8. Open a connection as root to a node in the cluster.
9. Deploy the new configuration:

```
dcli -C bdagetcclientconfig
```

[Figure 3-2](#) shows the Filesystem Trash Interval property in Cloudera Manager.

Figure 3–2 HDFS Property Settings in Cloudera Manager

The screenshot shows the Cloudera Manager interface for HDFS configuration. The top navigation bar includes 'Status', 'Instances', 'Commands', 'Configuration', 'Audits', 'Charts Library', and 'File Browser'. The 'Configuration' tab is active, showing a search bar with 'trash' and a 'Save Changes' button. Below the search bar, a green banner indicates '3 validation checks'. The configuration table lists two properties:

Category	Property	Value	Description
Gateway Default Group	Use Trash	<input checked="" type="checkbox"/> Reset to the default value: false HDFS Trash is enabled.	Move deleted files to the trash so that they can be recovered if necessary. This client side configuration takes effect only if the HDFS service-wide trash is disabled (NameNode Filesystem Trash Interval set to 0) and is ignored otherwise. The trash is not automatically emptied when enabled with this configuration.
NameNode Default Group	Filesystem Trash Interval	1 day(s) default value fs.trash.interval Trash checkpointing is on	Number of minutes between trash checkpoints. Also controls the number of minutes after which a trash checkpoint directory is deleted. To disable the trash feature, enter 0.

Disabling the Trash Facility

The trash facility on Oracle Big Data Appliance is enabled by default. You can change this configuration for a cluster. When the trash facility is disabled, deleted files and directories are not moved to the trash. They are not recoverable.

Completely Disabling the Trash Facility

The following procedure disables the trash facility for HDFS. When the trash facility is completely disabled, the client configuration is irrelevant.

To completely disable the trash facility:

1. Open Cloudera Manager. See ["Managing CDH Operations Using Cloudera Manager"](#) on page 2-3.
2. On the Home page under Status, click **hdfs**.
3. On the hdfs page, click the Configuration subtab, and then select **View and Edit**.
4. Search for or scroll down to the Filesystem Trash Interval property under NameNode Default Group. See [Figure 3–2](#).
5. Click the current value, and enter a value of 0 (zero) in the pop-up form.
6. Click **Save Changes**.
7. Expand the Actions menu at the top of the page and choose **Restart**.

Disabling the Trash Facility for Local HDFS Clients

All HDFS clients that are installed on Oracle Big Data Appliance are configured to use the trash facility. An **HDFS client** is any software that connects to HDFS to perform operations such as listing HDFS files, copying files to and from HDFS, and creating directories.

You can use Cloudera Manager to change the local client configuration setting, although the trash facility is still enabled.

Note: If you do not want any clients to use the trash, then you can completely disable the trash facility. See ["Completely Disabling the Trash Facility"](#) on page 3-13.

To disable the trash facility for local HDFS clients:

1. Open Cloudera Manager. See ["Managing CDH Operations Using Cloudera Manager"](#) on page 2-3.
2. On the Home page under Status, click **hdfs**.
3. On the hdfs page, click the **Configuration** subtab, and then select **View and Edit**.
4. Search for or scroll down to the Filesystem Trash Interval property under Gateway Default Group. See [Figure 3-2](#).
5. Search for or scroll down to the Use Trash property under Client Settings. See [Figure 3-2](#).
6. Deselect the Use Trash check box.
7. Click **Save Changes**. This setting is used to configure all new HDFS clients downloaded to Oracle Big Data Appliance.
8. Open a connection as **root** to a node in the cluster.
9. Deploy the new configuration:

```
dcli -C bdagetcclientconfig
```

Disabling the Trash Facility for a Remote HDFS Client

Remote HDFS clients are typically configured by downloading and installing a CDH client, as described in ["Providing Remote Client Access to CDH"](#) on page 3-2. Oracle SQL Connector for HDFS and Oracle R Advanced Analytics for Hadoop are examples of remote clients.

To disable the trash facility for a remote HDFS client:

1. Open a connection to the system where the CDH client is installed.
2. Open `/etc/hadoop/conf/hdfs-site.xml` in a text editor.
3. Set the trash interval to zero:

```
<property>
  <name>fs.trash.interval</name>
  <value>0</value>
</property>
```

4. Save the file.

Configuring Oracle Exadata Database Machine for Use with Oracle Big Data Appliance

This chapter provides information about optimizing communications between Oracle Exadata Database Machine and Oracle Big Data Appliance. It describes how you can configure Oracle Exadata Database Machine to use InfiniBand alone, or SDP over InfiniBand, to communicate with Oracle Big Data Appliance.

This chapter contains the following sections:

- [About Optimizing Communications](#)
- [Prerequisites for Optimizing Communications](#)
- [Specifying the InfiniBand Connections to Oracle Big Data Appliance](#)
- [Specifying the InfiniBand Connections to Oracle Exadata Database Machine](#)
- [Enabling SDP on Exadata Database Nodes](#)
- [Configuring a JDBC Client for SDP](#)
- [Creating an SDP Listener on the InfiniBand Network](#)

About Optimizing Communications

Oracle Exadata Database Machine and Oracle Big Data Appliance use Ethernet by default, although typically they are also connected by an InfiniBand network. Ethernet communications are much slower than InfiniBand. After you configure Oracle Exadata Database Machine to communicate using InfiniBand, it can obtain data from Oracle Big Data Appliance many times faster than before.

Moreover, client applications that run on Oracle Big Data Appliance and push the data to Oracle Database can use Sockets Direct Protocol (SDP) for an additional performance boost. SDP is a standard communication protocol for clustered server environments, providing an interface between the network interface card and the application. By using SDP, applications place most of the messaging burden upon the network interface card, which frees the CPU for other tasks. As a result, SDP decreases network latency and CPU utilization, and thereby improves performance.

About Applications that Pull Data Into Oracle Exadata Database Machine

Oracle SQL Connector for Hadoop Distributed File System (HDFS) is an example of an application that pulls data into Oracle Exadata Database Machine. The connector enables an Oracle external table to access data stored in either HDFS files or a Hive table.

The external table provide access to the HDFS data. You can use the external table for querying HDFS data or for loading it into an Oracle database table.

Oracle SQL Connector for HDFS functions as a Hadoop client running on the database servers in Oracle Exadata Database Machine.

If you use Oracle SQL Connector for HDFS or another tool that pulls the data into Oracle Exadata Database Machine, then for the best performance, you should configure the system to use InfiniBand. See ["Specifying the InfiniBand Connections to Oracle Big Data Appliance"](#) on page 4-2.

See Also : *Oracle Big Data Connectors User's Guide* for information about Oracle SQL Connector for HDFS

About Applications that Push Data Into Oracle Exadata Database Machine

Oracle Loader for Hadoop is an example of an application that pushes data into Oracle Exadata Database Machine. The connector is an efficient and high-performance loader for fast movement of data from a Hadoop cluster into a table in an Oracle database. You can use it to load data from Oracle Big Data Appliance to Oracle Exadata Database Machine.

Oracle Loader for Hadoop functions as a database client running on the Oracle Big Data Appliance. It must make database connections from Oracle Big Data Appliance to Oracle Exadata Database Machine over the InfiniBand network. Use of Sockets Direct Protocol (SDP) for these database connections further improves performance.

If you use Oracle Loader for Hadoop or another tool that pushes the data into Oracle Exadata Database Machine, then for the best performance, you should configure the system to use SDP over InfiniBand as described in this chapter.

See Also : *Oracle Big Data Connectors User's Guide* for information about Oracle Loader for Hadoop

Prerequisites for Optimizing Communications

Oracle Big Data Appliance and Oracle Exadata Database Machine racks must be cabled together using InfiniBand cables. The IP addresses must be unique across all racks and use the same subnet for the InfiniBand network.

See Also:

- *Oracle Big Data Appliance Owner's Guide* about multitrack cabling
- *Oracle Big Data Appliance Owner's Guide* about IP addresses and subnets

Specifying the InfiniBand Connections to Oracle Big Data Appliance

You can configure Oracle Exadata Database Machine to use the InfiniBand IP addresses of the Oracle Big Data Appliance servers. Otherwise, the default network is Ethernet. Use of the InfiniBand network improves the performance of all data transfers between Oracle Big Data Appliance and Oracle Exadata Database Machine.

To identify the Oracle Big Data Appliance InfiniBand IP addresses:

1. If you have not done so already, install a CDH client on Oracle Exadata Database Machine. See ["Providing Remote Client Access to CDH"](#) on page 3-2.

2. Obtain a list of private host names and InfiniBand IP addresses for all Oracle Big Data Appliance servers.

An Oracle Big Data Appliance rack can have 6, 12, or 18 servers.

3. Log in to Oracle Exadata Database Machine with `root` privileges.
4. Edit `/etc/hosts` on Oracle Exadata Database Machine and add the Oracle Big Data Appliance host names and InfiniBand IP addresses. The following example shows the sequential IP numbering:

```
192.168.8.1      bda1node01.example.com  bda1node01
192.168.8.2      bda1node02.example.com  bda1node02
192.168.8.3      bda1node03.example.com  bda1node03
192.168.8.4      bda1node04.example.com  bda1node04
192.168.8.5      bda1node05.example.com  bda1node05
192.168.8.6      bda1node06.example.com  bda1node06
```

5. Check `/etc/nsswitch.conf` for a line like the following:

```
hosts:          files dns
```

Ensure that the line does not reverse the order (`dns files`); if it does, your additions to `/etc/hosts` will not be used. Edit the file if necessary.

6. Ping all Oracle Big Data Appliance servers. Ensure that ping completes and shows the InfiniBand IP addresses.

```
# ping bda1node01.example.com
PING bda1node01.example.com (192.168.8.1) 56(84) bytes of data.
64 bytes from bda1node01.example.com (192.168.8.1): icmp_seq=1 ttl=50 time=20.2
ms
.
.
.
```

7. Run CDH locally on Oracle Exadata Database Machine and test HDFS functionality by uploading a large file to an Oracle Big Data Appliance server. Check that your network monitoring tools (such as `sar`) show I/O activity on the InfiniBand devices.

To upload a file, use syntax like the following, which copies `localfile.dat` to the HDFS `testdir` directory on node05 of Oracle Big Data Appliance:

```
hadoop fs -put localfile.dat hdfs://bda1node05.example.com/testdir/
```

Specifying the InfiniBand Connections to Oracle Exadata Database Machine

You can also configure Oracle Big Data Appliance to use the InfiniBand IP addresses of the Oracle Exadata Database Machine servers. This configuration supports applications on Oracle Big Data Appliance that must connect to Oracle Exadata Database Machine.

To identify the Oracle Exadata Database Machine InfiniBand IP addresses:

1. Obtain a list of private host names and InfiniBand IP addresses for all Oracle Exadata Database Machine servers.
2. Log in to Oracle Big Data Appliance with `root` privileges.

3. Edit `/etc/hosts` on Oracle Big Data Appliance and add the Oracle Exadata Database Machine host names and InfiniBand IP addresses.
4. Check `/etc/nsswitch.conf` for a line like the following:

```
hosts:      files dns
```

Ensure that the line does not reverse the order (`dns files`); if it does, your additions to `/etc/hosts` will not be used. Edit the file if necessary.
5. Restart the `dnsmasq` service:

```
# service dnsmasq restart
```
6. Ping all Oracle Exadata Database Machine servers. Ensure that ping completes and shows the InfiniBand IP addresses.
7. Test the connection by downloading a large file to an Oracle Exadata Database Machine server. Check that your network monitoring tools (such as `sar`) show I/O activity on the InfiniBand devices.

To download a file, use syntax like the following, which copies a file named `mydata.json` to the `dm01cel08` storage server:

```
$ scp mydata.json oracle@dm01cel08-priv.example.com:mybigdata.json
oracle@dm01cel08-priv.example.com's password: password
```

Enabling SDP on Exadata Database Nodes

The following procedure describes how to enable SDP on database nodes in an Oracle Exadata Database Machine running Oracle Linux. SDP improves the performance of client applications that run on Oracle Big Data Appliance and push data to Oracle Exadata Database Machine.

To enable SDP on Oracle Exadata Database Machine:

1. Open `/etc/infiniband/openib.conf` file in a text editor, and add the following line:

```
set: SDP_LOAD=yes
```
2. Save these changes and close the file.
3. To enable both SDP and TCP, open `/etc/ofed/libsdp.conf` in a text editor, and add the `use both` rule:

```
use both server * :
use both client * :
```
4. Save these changes and close the file.
5. Open `/etc/modprobe.conf` file in a text editor, and add this setting:

```
options ib_sdp sdp_zcopy_thresh=0 recv_poll=0
```
6. Save these changes and close the file.
7. Replicate these changes across all database nodes in the Oracle Exadata Database Machine rack.
8. Restart all database nodes for the changes to take effect.
9. If you have multiple Oracle Exadata Database Machine racks, then repeat these steps on all of them.

Configuring a JDBC Client for SDP

The following procedure explains how to configure a JDBC client to use SDP.

To enable SDP support for JDBC:

1. Configure the database to support InfiniBand, as described in the *Oracle Database Net Services Administrator's Guide*. Ensure that you set the protocol to SDP.
2. Set the LD_PRELOAD environment variable to libsdp.so before starting the Java virtual machine. This example uses the Bash shell:

```
export LD_PRELOAD="libsdp.so"
```

Creating an SDP Listener on the InfiniBand Network

To add a listener for the Oracle Big Data Appliance connections coming in on the InfiniBand network, first add a network resource for the InfiniBand network with virtual IP addresses.

Note: This example lists two nodes for an Oracle Exadata Database Machine quarter rack. If you have an Oracle Exadata Database Machine half or full rack, you must repeat node-specific lines for each node in the cluster.

1. Edit /etc/hosts on each node in the Exadata rack to add the virtual IP addresses for the InfiniBand network. Make sure that these IP addresses are not in use. For example:

```
# Added for Listener over IB
192.168.10.21 dm01db01-ibvip.example.com dm01db01-ibvip
192.168.10.22 dm01db02-ibvip.example.com dm01db02-ibvip
```

2. As the root user, create a network resource on one database node for the InfiniBand network. For example:

```
# /u01/app/grid/product/12.1.0.1/bin/srvctl add network -k 2 -S
192.168.10.0/255.255.255.0/bondib0
```

3. Verify that the network was added correctly with a command like the following examples:

```
# /u01/app/grid/product/12.1.0.1/bin/crsctl stat res -t | grep net
ora.net1.network
ora.net2.network -- Output indicating new Network resource
```

or

```
# /u01/app/grid/product/12.1.0.1/bin/srvctl config network -k 2
Network exists: 2/192.168.10.0/255.255.255.0/bondib0, type static -- Output
indicating Network resource on the 192.168.10.0 subnet
```

4. Add the virtual IP addresses on the network created in Step 2, for each node in the cluster. For example:

```
# srvctl add vip -n dm01db01 -A dm01db01-ibvip/255.255.255.0/bondib0 -k 2
#
# srvctl add vip -n dm01db02 -A dm01db02-ibvip/255.255.255.0/bondib0 -k 2
```

5. As the oracle user who owns Grid Infrastructure Home, add a listener for the virtual IP addresses created in Step 4.

```
# srvctl add listener -l LISTENER_IB -k 2 -p TCP:1522,/SDP:1522
```

6. For each database that will accept connections from the middle tier, modify the `listener_networks` init parameter to allow load balancing and failover across multiple networks (Ethernet and InfiniBand). You can either enter the full TNSNAMES syntax in the initialization parameter or create entries in `tnsnames.ora` in the `$ORACLE_HOME/network/admin` directory. The TNSNAMES.ORA entries must exist in GRID_HOME. The following example first updates `tnsnames.ora`.

Complete this step on each node in the cluster with the correct IP addresses for that node. LISTENER_IBREMOTE should list all other nodes that are in the cluster. DBM_IB should list all nodes in the cluster.

Note: The database instance reads the TNSNAMES only on startup. Thus, if you modify an entry that is referred to by any `init.ora` parameter (`LISTENER_NETWORKS`), then you must either restart the instance or issue an `ALTER SYSTEM SET LISTENER_NETWORKS` command for the modifications to take affect by the instance.

```
DBM =
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP) (HOST = dm01-scan) (PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = dbm)
  )
)

DBM_IB =
(DESCRIPTION =
  (LOAD_BALANCE=on)
  (ADDRESS = (PROTOCOL = TCP) (HOST = dm01db01-ibvip) (PORT = 1522))
  (ADDRESS = (PROTOCOL = TCP) (HOST = dm01db02-ibvip) (PORT = 1522))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = dbm)
  )
)

LISTENER_IBREMOTE =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = dm01db02-ibvip.mycompany.com) (PORT = 1522))
  )
)

LISTENER_IBLOCAL =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = dm01db01-ibvip.mycompany.com) (PORT = 1522))
    (ADDRESS = (PROTOCOL = SDP) (HOST = dm01db01-ibvip.mycompany.com) (PORT = 1523))
  )
)

LISTENER_IPLOCAL =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP) (HOST = dm0101-vip.mycompany.com) (PORT = 1521))
  )
)
```

```

LISTENER_IPREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP) (HOST = dm01-scan.mycompany.com) (PORT = 1521))
))

```

7. Connect to the database instance as sysdba.
8. Modify the `listener_networks` init parameter by using the SQL `ALTER SYSTEM` command:

```

SQL> alter system set listener_networks=
      '((NAME=network2) (LOCAL_LISTENER=LISTENER_IBLOCAL)
        (REMOTE_LISTENER=LISTENER_IBREMOTE))',
      '((NAME=network1) (LOCAL_LISTENER=LISTENER_IPLOCAL)
        (REMOTE_LISTENER=LISTENER_IPREMOTE))' scope=both;

```

9. On the Linux command line, use the `srvctl` command to restart `LISTENER_IB` to implement the modification in Step 7:

```

# srvctl stop listener -l LISTENER_IB
# srvctl start listener -l LISTENER_IB

```

Optimizing MapReduce Jobs Using Perfect Balance

This chapter describes how you can shorten the run time of some MapReduce jobs by using Perfect Balance. It contains the following sections:

- [What is Perfect Balance?](#)
- [Application Requirements](#)
- [Getting Started with Perfect Balance](#)
- [Analyzing a Job's Reducer Load](#)
- [About Configuring Perfect Balance](#)
- [Running a Balanced MapReduce Job Using Perfect Balance](#)
- [About Perfect Balance Reports](#)
- [About Chopping](#)
- [Troubleshooting Jobs Running with Perfect Balance](#)
- [Using the Perfect Balance API](#)
- [About the Perfect Balance Examples](#)
- [Perfect Balance Configuration Property Reference](#)

What is Perfect Balance?

The Perfect Balance feature of Oracle Big Data Appliance distributes the reducer load in a MapReduce application so that each reduce task does approximately the same amount of work. While the default Hadoop method of distributing the reduce load is appropriate for many jobs, it does not distribute the load evenly for jobs with significant data skew.

Data skew is an imbalance in the load assigned to different reduce tasks. The **load** is a function of:

- The number of keys assigned to a reducer.
- The number of records and the number of bytes in the values per key.

The total run time for a job is extended, to varying degrees, by the time that the reducer with the greatest load takes to finish. In jobs with a skewed load, some reducers complete the job quickly, while others take much longer. Perfect Balance can significantly shorten the total run time by distributing the load evenly, enabling all reducers to finish at about the same time.

Your MapReduce job can be written using either the `mapred` or `mapreduce` APIs; Perfect Balance supports both of them.

Perfect Balance was tested on MapReduce 1 (MRv1) and YARN clusters. YARN is the default installation on Oracle Big Data Appliance 3.0 and later releases. See ["Switching From YARN to MapReduce 1"](#) on page 2-22 for the Perfect Balance configuration settings when switching Oracle Big Data Appliance from YARN to MRv1.

About Balancing Jobs Across Map and Reduce Tasks

A typical Hadoop job has map and reduce tasks. Hadoop distributes the *mapper* workload uniformly across Hadoop Distributed File System (HDFS) and across map tasks, while preserving the data locality. In this way, it reduces skew in the mappers.

Hadoop also hashes the map-output keys uniformly across all *reducers*. This strategy works well when there are many more keys than reducers, and each key represents a very small portion of the workload. However, it is not effective when the mapper output is concentrated into a small number of keys. Hashing these keys results in skew and does not work in applications like sorting, which require range partitioning.

Perfect Balance distributes the load evenly across reducers by first sampling the data, optionally chopping large keys into two or more smaller keys, and using a load-aware partitioning strategy to assign keys to reduce tasks.

Ways to Use Perfect Balance Features

You can choose from two methods of running Perfect Balance features:

- **Perfect Balance:** You can run a job without changing your application code by properly configuring Perfect Balance. This is the preferred method and appropriate for most jobs.

This method is described in ["Running a Balanced MapReduce Job Using Perfect Balance"](#) on page 5-10, and is the primary focus of this chapter.

- **Perfect Balance API:** You can add the code to run Perfect Balance to your application code. Use this method when your application setup code must run before using Perfect Balance. Otherwise, you should use the previous method, which requires no change to your code.

This method is described in ["Using the Perfect Balance API"](#) on page 5-15.

Perfect Balance Components

Perfect Balance has these components:

- **Job Analyzer:** Gathers and reports statistics about the MapReduce job so that you can determine whether to use Perfect Balance.
- **Counting Reducer:** Provides additional statistics to the Job Analyzer to help gauge the effectiveness of Perfect Balance.
- **Load Balancer:** Runs before the MapReduce job to generate a static partition plan, and reconfigures the job to use the plan. The balancer includes a user-configurable, progressive sampler that stops sampling the data as soon as it can generate a good partitioning plan.

Application Requirements

To use Perfect Balance successfully, your application must meet the following requirements:

- The job is distributive, so that splitting a group of records associated with a reduce key does not produce incorrect results for the application.

To balance a load, Perfect Balance subpartitions the values of large reduce keys and sends each subpartition to a different reducer. This distribution contrasts with the standard Hadoop practice of sending all values for a single reduce key to the same reducer. Your application must be able to handle output from the reducers that is not fully aggregated, so that it does not produce incorrect results.

This partitioning of values is called **chopping**. Applications that support chopping have **distributive reduce functions**. See ["About Chopping"](#) on page 5-13.

If your application is not distributive, then you can still run Perfect Balance after disabling the key-splitting feature. The job still benefits from using Perfect Balance, but the load is not as evenly balanced as it is when key splitting is in effect. See the [oracle.hadoop.balancer.keyLoad.minChopBytes](#) configuration property to disable key splitting.

- This release does not support combiners. Perfect Balance detects the presence of combiners and does not balance when they are present.

Getting Started with Perfect Balance

Take the following steps to use Perfect Balance:

1. Ensure that your application meets the requirements listed in ["Application Requirements"](#) on page 5-3.
2. Log in to the server where you will submit the job.
3. Run the examples provided with Perfect Balance to become familiar with the product. All examples shown in this chapter are based on the shipped examples and use the same data set. See ["About the Perfect Balance Examples"](#) on page 5-17.
4. Set the following variables using the Bash `export` command:
 - `BALANCER_HOME`: Set to the Perfect Balance installation directory, such as `/opt/oracle/orabalancer-2.0.0-h2` on Oracle Big Data Appliance (optional). The examples in this chapter use this variable, and you can also define it for your convenience. Perfect Balance does not require `BALANCER_HOME`.
 - `HADOOP_CLASSPATH`: Add `${BALANCER_HOME}/jlib/orabalancer-2.0.0.jar` and `${BALANCER_HOME}/jlib/commons-math-2.2.jar` to the existing value. Also add the JAR files for your application.

To enable Perfect Balance, add `${BALANCER_HOME}/jlib/orabalancerclient-2.0.0.jar`. (API mode does not use this JAR file.)
 - `HADOOP_USER_CLASSPATH_FIRST`: Set to `true` to enable Perfect Balance; API mode does not need this variable.

You do not need to set these variables unless you are using Perfect Balance.

5. Run Job Analyzer without the balancer and use the generated report to decide whether the job is a good candidate for using Perfect Balance.

See ["Analyzing a Job's Reducer Load"](#) on page 5-4.

6. Decide which configuration properties to set. Create a configuration file or enter the settings individually in the `hadoop` command.
See ["About Configuring Perfect Balance"](#) on page 5-9.
7. Run the job using Perfect Balance.
See ["Running a Balanced MapReduce Job Using Perfect Balance"](#) on page 5-10.
8. Use the Job Analyzer report to evaluate the effectiveness of using Perfect Balance.
See ["Reading the Job Analyzer Report"](#) on page 5-8.
9. Modify the job configuration properties as desired before rerunning the job with Perfect Balance. See ["About Configuring Perfect Balance"](#) on page 5-9.

Analyzing a Job's Reducer Load

Job Analyzer is a component of Perfect Balance that identifies imbalances in a load, and how effective Perfect Balance is in correcting the imbalance when actually running the job. This section contains the following topics:

- [About Job Analyzer](#)
- [Running Job Analyzer as a Standalone Utility](#)
- [Running Job Analyzer Using Perfect Balance](#)
- [Reading the Job Analyzer Report](#)

About Job Analyzer

You can use Job Analyzer to decide whether a job is a candidate for load balancing with Perfect Balance. Job Analyzer uses the output logs of a MapReduce job to generate a simple report with statistics like the elapsed time and the load for each reduce task. By default, it uses the standard Hadoop counters displayed by the JobTracker user interface, but organizes the data to emphasize the relative performance and load of the reduce tasks, so that you can more easily interpret the results.

If the report shows that the data is skewed (that is, the reducers processed very different loads and the run times varied widely), then the application is a good candidate for Perfect Balance.

Methods of Running Job Analyzer

You can choose between two methods of running Job Analyzer:

- **As a standalone utility:** Job Analyzer runs against existing job output logs. This is a good choice when you want to analyze a job that previously ran.
- **While using Perfect Balance:** Job Analyzer runs against the output logs for the current job running with Perfect Balance. This is a good choice when you want to analyze the current job.

Running Job Analyzer as a Standalone Utility

As a standalone utility, Job Analyzer provides a quick way to analyze the reduce load of a previously run job.

To run Job Analyzer as a standalone utility:

1. Log in to the server where you will run Job Analyzer.

2. Locate the output logs from the job to analyze:

- **YARN clusters:** Set `oracle.hadoop.balancer.application_id` to the job ID of the job you want to analyze. YARN is the default on Oracle Big Data Appliance.

You can obtain the job ID from the YARN Resource Manager web interface. Click the application ID of a job, and then click **Tracking URL**. The job ID typically begins with "job_".

Alternately, if you already ran Perfect Balance or Job Analyzer on this job, you can read the job ID from the `application_id` file generated by Perfect Balance in its report directory (`outdir/_balancer` by default).

- **MRv1 clusters:** Set `mapred.output.dir` to the output directory of the job you want to analyze.

3. Run the Job Analyzer utility as described in "Job Analyzer Utility Syntax" on page 5-5.

4. View the Job Analyzer report in a browser.

Job Analyzer Utility Example

Example 5-1 runs a script that sets the required variables, uses the MapReduce job logs for a job with an application ID of `job_1396563311211_0947`, and creates the report in the default location. It then copies the HTML version of the report from HDFS to the `/home/jdoe` local directory and opens the report in a browser.

If you want to run this example in YARN, then replace the application ID with the application ID of the job. The application ID of the job looks like this example: `job_1396563311211_0947`.

If you use MRv1 instead of YARN, then set `mapred.output.dir` instead of `oracle.hadoop.balancer.application_id`.

Example 5-1 Running the Job Analyzer Utility

```
$ cat runja.sh

BALANCER_HOME=/opt/oracle/orabalancer-2.0.0-h2
export HADOOP_CLASSPATH=${BALANCER_HOME}/jlib/orabalancer-2.0.0.jar:${BALANCER_
HOME}/jlib/commons-math-2.2.jar:$HADOOP_CLASSPATH
export HADOOP_USER_CLASSPATH_FIRST=true

# Command on YARN cluster
hadoop jar orabalancer.jar oracle.hadoop.balancer.tools.JobAnalyzer \
-D oracle.hadoop.balancer.application_id=job_1396563311211_0947

$ sh ./runja.sh
$
$ hadoop fs -get jdoe_nobal_outdir/_balancer/jobanalyzer-report.html /home/jdoe
$ cd /home/jdoe
$ firefox jobanalyzer-report.html
```

Job Analyzer Utility Syntax

The following is the syntax to run the Job Analyzer utility:

For YARN:

```
hadoop jar orabalancer.jar oracle.hadoop.balancer.tools.JobAnalyzer \
-D oracle.hadoop.balancer.application_id=job_number
```

For MRv1:

```
hadoop jar ${BALANCER_HOME}/jlib/orabalancer-2.0.0
oracle.hadoop.balancer.tools.JobAnalyzer \
-D mapred.output.dir=job_output_dir \
[ja_report_path]
```

job_number

The application ID previously assigned to the job. YARN only.

job_output_dir

An HDFS directory where the job files are stored from previously executing the application. MRv1 only.

ja_report_path

An HDFS directory where Job Analyzer creates its report (optional). The default directory is *job_output_dir/_balancer*. YARN only.

Running Job Analyzer Using Perfect Balance

When you run a job using Perfect Balance, you can configure it to run Job Analyzer automatically. This section contains the following topics:

- [Running Job Analyzer Using Perfect Balance](#)
- [Collecting Additional Metrics](#)

Running Job Analyzer Using Perfect Balance

Follow these steps to run Job Analyzer using Perfect Balance:

1. Log in to the server where you will submit the job that uses Perfect Balance.
2. Set up Perfect Balance by taking the steps in "[Getting Started with Perfect Balance](#)" on page 5-3.
3. To enable Job Analyzer, set the `oracle.hadoop.balancer.autoAnalyze` configuration property to one of these values:
 - `BASIC_REPORT`: Enables Job Analyzer. If you set `oracle.hadoop.balancer.autoBalance` to true, then Perfect Balance automatically sets `oracle.hadoop.balancer.autoAnalyze` to `BASIC_REPORT`.
 - `REDUCER_REPORT`: Configures Job Analyzer to collect additional load statistics. See "[Collecting Additional Metrics](#)" on page 5-7.
4. Decide which additional configuration properties to set, if any.
See "[Perfect Balance Configuration Property Reference](#)" on page 5-18.
5. Run the job.

Example 5-2 runs a script that sets the required variables, uses Perfect Balance to run a job with Job Analyzer and without load balancing, and creates the report in the default location. It then copies the HTML version of the report from HDFS to the `/home/jdoe` local directory and opens the report in a browser. The output includes warnings, which you can ignore.

Example 5-2 Running Job Analyzer with Perfect Balance

```
$ cat ja_nobalance.sh
```

```
# set up perfect balance
BALANCER_HOME=/opt/oracle/orabalancer-2.0.0-h2
export HADOOP_CLASSPATH=${BALANCER_
HOME}/jlib/orabalancerclient-2.0.0.jar:${BALANCER_
HOME}/jlib/orabalancer-2.0.0.jar:${BALANCER_
HOME}/jlib/commons-math-2.2.jar:${HADOOP_CLASSPATH}
export HADOOP_USER_CLASSPATH_FIRST=true

# run the job
hadoop jar application_jarfile.jar ApplicationClass \
-D application_config_property \
-D mapreduce.input.fileinputformat.inputdir=jdoe_application/input \
-D mapreduce.output.fileoutputformat.outputdir=jdoe_nobal_outdir \
-D mapreduce.job.name=nobal \
-D mapreduce.job.reduces=10 \
-D oracle.hadoop.balancer.autoBalance=false \
-D oracle.hadoop.balancer.autoAnalyze=REDUCER_REPORT \
-conf application_config_file.xml

$ sh ja_nobalance.sh
14/04/14 14:52:42 INFO input.FileInputFormat: Total input paths to process : 5
14/04/14 14:52:42 INFO mapreduce.JobSubmitter: number of splits:5
14/04/14 14:52:42 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_
1397066986369_3478
14/04/14 14:52:43 INFO impl.YarnClientImpl: Submitted application application_
1397066986369_3478
.
.
.
File Input Format Counters
Bytes Read=112652976
File Output Format Counters
Bytes Written=384974202

$ hadoop fs -get jdoe_nobal_outdir/_balancer/jobanalyzer-report.html /home/jdoe
$ cd /home/jdoe
$ firefox jobanalyzer-report.html
```

Collecting Additional Metrics

The Job Analyzer report includes the load metrics for each key, if you set the `oracle.hadoop.balancer.autoAnalyze` property to `REDUCER_REPORT`. This additional information provides a more detailed picture of the load for each reducer, with metrics that are not available in the standard Hadoop counters.

The Job Analyzer report also compares its predicted load with the actual load. The difference between these values measures how effective Perfect Balance was in balancing the job.

Job Analyzer might recommend key load coefficients for the Perfect Balance key load model, based on its analysis of the job load. To use these recommended coefficients when running a job with Perfect Balance, set the `oracle.hadoop.balancer.linearKeyLoad.feedbackDir` property to the directory containing the Job Analyzer report of a previously analyzed run of the job.

If the report contains recommended coefficients, then Perfect Balance automatically uses them. If Job Analyzer encounters an error while collecting the additional metrics, then the report does not contain the additional metrics.

Use the `feedbackDir` property when you do not know the values of the load model coefficients for a job, but you have the Job Analyzer output from a previous run of the

job. Then you can set the value of `feedbackDir` to the directory where that output is stored. The values recommended from those files typically perform better than the Perfect Balance default values, because the recommended values are based on an analysis of your job's load.

Alternately, if you already know good values of the load model coefficients for your job, you can set the load model properties:

- `oracle.hadoop.balancer.linearKeyLoad.byteWeight`
- `oracle.hadoop.balancer.linearKeyLoad.keyWeight`
- `oracle.hadoop.balancer.linearKeyLoad.rowWeight`

Running the job with these coefficients results in a more balanced job.

Reading the Job Analyzer Report

Job Analyzer writes its report in two formats: HTML for you, and XML for Perfect Balance. You can open the report in a browser, either directly in HDFS or after copying it to the local file system

To open a Job Analyzer report in HDFS in a browser:

1. Open the HDFS web interface on port 50070 of a NameNode node (node01 or node02), using a URL like the following:

```
http://bdainode01.example.com:50070
```

2. From the Utilities menu, choose **Browse the File System**.
3. Navigate to the `job_output_dir/_balancer` directory.

To open a Job Analyzer report in the local file system in a browser:

1. Copy the report from HDFS to the local file system:

```
$ hadoop fs -get job_output_dir/_balancer/jobanalyzer-report.html /home/jdoe
```

2. Switch to the local directory:

```
$ cd /home/jdoe
```

3. Open the file in a browser:

```
$ firefox jobanalyzer-report.html
```

When inspecting the Job Analyzer report, look for indicators of skew such as:

- The execution time of some reducers is longer than others.
- Some reducers process more records or bytes than others.
- Some map output keys have more records than others.
- Some map output records have more bytes than others.

[Figure 5-1](#) shows the beginning of the analyzer report for the inverted index (`invindx`) example. It displays the key load coefficient recommendations, because this job ran with the appropriate configuration settings. See ["Collecting Additional Metrics"](#) on page 5-7.

The task IDs are links to tables that show the analysis of specific tasks, enabling you to drill down for more details from the first, summary table.

This example uses an extremely small data set, but notice the differences between tasks 7 and 8: The input records range from 3% to 29%, and their corresponding elapsed times range from 5 to 15 seconds. This variation indicates skew.

Figure 5–1 Job Analyzer Report for Unbalanced Inverted Index Job

Job Information

Job Name	invindx
Job Id	job_1396563311211_0959
Start Time	2014-04-08 06:51:52
Finish Time	2014-04-08 06:52:38
Map Time	00:00:25
Reduce Time	00:00:16
Elapsed Time	00:00:46

Reduce Tasks

Task ID	Time				Input								Output				%Load	
	Start	Finish	Elapsed	CPU		Shuffle Bytes		Keys		Records		ValueBytes		Records		Bytes		Observed
				time	%	count	%	count	%	count	%	count	%	count	%	count	%	
0	06:52:22	06:52:31	00:00:09	00:00:11	11	20,987,702	12	13	13	2,300,118	12	43,702,242	12	1,698,171	12	48,022,226	12	12
1	06:52:22	06:52:30	00:00:07	00:00:09	9	15,630,571	9	12	12	1,438,876	7	27,338,644	7	1,309,838	9	36,424,492	9	7
2	06:52:22	06:52:30	00:00:07	00:00:09	9	14,682,583	8	7	7	1,503,328	8	28,563,232	8	1,220,301	9	33,206,894	9	8
3	06:52:22	06:52:30	00:00:07	00:00:09	9	14,977,081	9	10	10	1,450,603	7	27,561,457	7	1,253,630	9	33,638,218	9	7
4	06:52:22	06:52:33	00:00:11	00:00:11	11	21,175,220	12	11	11	2,073,885	10	39,403,815	10	1,768,698	13	51,191,915	13	10
5	06:52:22	06:52:32	00:00:09	00:00:11	11	19,664,004	11	10	10	2,580,271	13	49,025,149	13	1,497,856	11	40,709,092	11	13
6	06:52:23	06:52:29	00:00:06	00:00:07	7	8,867,408	5	9	9	775,690	4	14,738,110	4	740,944	5	20,503,323	5	4
7	06:52:23	06:52:29	00:00:05	00:00:06	6	5,912,478	3	6	6	531,247	3	10,093,693	3	494,798	4	13,336,988	3	3
8	06:52:23	06:52:38	00:00:15	00:00:18	18	35,890,632	21	11	11	5,824,950	29	110,674,050	29	2,527,641	18	69,087,902	18	29
9	06:52:23	06:52:31	00:00:07	00:00:09	9	16,252,928	9	11	11	1,521,032	8	28,899,608	8	1,359,917	10	38,889,633	10	8
Total	-	-	-	00:01:43	-	174,040,607	-	100	-	20,000,000	-	380,000,000	-	13,871,794	-	385,010,683	-	-

About Configuring Perfect Balance

Perfect Balance uses the standard Hadoop methods of specifying configuration properties in the command line. You can use the `-conf` option to identify a configuration file, or the `-D` option to specify individual properties. All Perfect Balance configuration properties have default values, and so setting them is optional.

"[Perfect Balance Configuration Property Reference](#)" on page 5-18 lists the configuration properties in alphabetical order with a full description. The following are functional groups of properties.

Perfect Balance Basic Properties

- `oracle.hadoop.balancer.autoAnalyze`
- `oracle.hadoop.balancer.autoBalance`

Job Analyzer Properties

- `oracle.hadoop.balancer.application_id`
- `oracle.hadoop.balancer.tools.jobConfPath`
- `oracle.hadoop.balancer.tools.jobHistoryPath`
- `oracle.hadoop.balancer.tools.writeKeyBytes`

Key Chopping Properties

- `oracle.hadoop.balancer.enableSorting`
- `oracle.hadoop.balancer.keyLoad.minChopBytes`

Load Balancing Properties

- `oracle.hadoop.balancer.confidence`
- `oracle.hadoop.balancer.maxLoadFactor`
- `oracle.hadoop.balancer.maxSamplesPct`
- `oracle.hadoop.balancer.minSplits`

Load Model Properties

- `oracle.hadoop.balancer.linearKeyLoad.feedbackDir`
- `oracle.hadoop.balancer.linearKeyLoad.byteWeight`
- `oracle.hadoop.balancer.linearKeyLoad.keyWeight`
- `oracle.hadoop.balancer.linearKeyLoad.rowWeight`

MapReduce-Related Properties

- `oracle.hadoop.balancer.useMapreduceApi`
- `oracle.hadoop.balancer.inputFormat.mapred.map.tasks`
- `oracle.hadoop.balancer.inputFormat.mapred.max.split.size`

Partition Report Properties

- `oracle.hadoop.balancer.report.override`
- `oracle.hadoop.balancer.reportPath`
- `oracle.hadoop.balancer.tmpDir`

Sampler Properties

- `oracle.hadoop.balancer.minSplits`
- `oracle.hadoop.balancer.numThreads`
- `oracle.hadoop.balancer.runMode`
- `oracle.hadoop.balancer.useClusterStats`

Running a Balanced MapReduce Job Using Perfect Balance

Perfect Balance does not require you to make any changes to your application code. It works by automatically running Perfect Balance for your job when you submit it to Hadoop for execution.

To run a job with Perfect Balance:

1. Log in to the server where you will submit the job.
2. Set up Perfect Balance by following the steps in ["Getting Started with Perfect Balance"](#) on page 5-3.
3. Configure the job with these Perfect Balance properties:
 - To enable balancing, set `oracle.hadoop.balancer.autoBalance` to true. This setting also runs Job Analyzer. Load balancing is not enabled by default.

- To allow Job Analyzer to collect additional metrics, set `oracle.hadoop.balancer.autoAnalyze` to `REDUCER_REPORT`.
See ["Collecting Additional Metrics"](#) on page 5-7.
- Decide which additional configuration properties to set, if any.
See ["About Configuring Perfect Balance"](#) on page 5-9.

4. Run your job as usual, using the following syntax:

```
bin/hadoop jar application_jarfile.jar ApplicationClass \
-D application_config_property \
-D oracle.hadoop.balancer.autoBalance=true \
-D other_perfect_balance_config_property \
-conf application_config_file.xml \
-conf perfect_balance_config_file.xml
```

You do not need to make any code changes to your application. You can provide Perfect Balance configuration properties either on the command line or in a configuration file. You can also combine Perfect Balance properties and MapReduce properties in the same configuration file. ["About Configuring Perfect Balance"](#) on page 5-9.

[Example 5-3](#) runs a script named `pb_balance.sh`, which sets up Perfect Balance for a job, and then runs the job. The key load metric properties are set to the values recommended in the Job Analyzer report shown in [Figure 5-1](#).

Example 5-3 Running a Job Using Perfect Balance

```
$ cat pb_balance.sh
```

```
#setup perfect balance as described in Getting Started with Perfect Balance
BALANCER_HOME=/opt/oracle/orabalancer-2.0.0-h2
export HADOOP_CLASSPATH=${BALANCER_HOME}/jlib/orabalancer-2.0.0.jar:${BALANCER_
HOME}/jlib/orabalancer.jar:${BALANCER_HOME}/jlib/commons-math-2.2.jar:$HADOOP_
CLASSPATH
export HADOOP_USER_CLASSPATH_FIRST=true

# setup optional properties like java heap size and garbage collector
export HADOOP_CLIENT_OPTS="-XX:+UseConcMarkSweepGC -Xmx512M ${HADOOP_CLIENT_OPTS}"

# run the job with balancing and job analyzer enabled
hadoop jar application_jarfile.jarApplicationClass
-D application_config_property \
-D mapreduce.input.fileinputformat.inputdir=jdoe_application/input \
-D mapreduce.output.fileoutputformat.outputdir=jdoe_outdir \
-D mapreduce.job.name="autoinvoke" \
-D mapreduce.job.reduces=10 \
-D oracle.hadoop.balancer.autoBalance=true \
-D oracle.hadoop.balancer.autoAnalyze=REDUCER_REPORT \
-D oracle.hadoop.balancer.linearKeyLoad.keyWeight=93.98 \
-D oracle.hadoop.balancer.linearKeyLoad.rowWeight=0.001126 \
-D oracle.hadoop.balancer.linearKeyLoad.byteWeight=0.0 \
-conf application_config_file.xml
```

```
$ sh ./pb_balance.sh
14/04/14 14:59:42 INFO balancer.Balancer: Creating balancer
14/04/14 14:59:42 INFO balancer.Balancer: Starting Balancer
14/04/14 14:59:43 INFO input.FileInputFormat: Total input paths to process : 5
14/04/14 14:59:46 INFO balancer.Balancer: Balancer completed
```

```

14/04/14 14:59:47 INFO input.FileInputFormat: Total input paths to process : 5
14/04/14 14:59:47 INFO mapreduce.JobSubmitter: number of splits:5
14/04/14 14:59:47 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_
1397066986369_3500
14/04/14 14:59:47 INFO impl.YarnClientImpl: Submitted application application_
1397066986369_3500
14/04/14 14:59:47 INFO mapreduce.Job: The url to track the job:
.
.
.
Map-Reduce Framework
Map input records=1000000
Map output records=20000000
Map output bytes=872652976
Map output materialized bytes=175650573
Input split bytes=580
Combine input records=0
Combine output records=0
Reduce input groups=106
Reduce shuffle bytes=175650573
Reduce input records=20000000
Reduce output records=13871794
Spilled Records=60000000
Shuffled Maps =50
Failed Shuffles=0
Merged Map outputs=50
GC time elapsed (ms)=1573
CPU time spent (ms)=242850
Physical memory (bytes) snapshot=6789033984
Virtual memory (bytes) snapshot=24548044800
Total committed heap usage (bytes)=11921457152
Shuffle Errors
BAD_ID=0
CONNECTION=0
IO_ERROR=0
WRONG_LENGTH=0
WRONG_MAP=0
WRONG_REDUCE=0
File Input Format Counters
Bytes Read=112652976
File Output Format Counters
Bytes Written=384974202

```

About Perfect Balance Reports

Perfect Balance generates these reports when it runs a job:

- **Job Analyzer report:** Contains various indicators about the distribution of the load in a job. The report is saved in HTML for you, and XML for Perfect Balance to use. The report is always named `jobanalyzer-report.html` and `-.xml`. See ["Reading the Job Analyzer Report"](#) on page 5-8.
- **Partition report:** Identifies the keys that are assigned to the various mappers. This report is saved in XML for Perfect Balance to use; it does not contain information of use to you. The report is named `${job_output_dir}/_balancer/orabalancer_report.xml`. It is only generated for balanced jobs.
- **Reduce key metric reports:** Perfect Balance generates a report for each file partition, when the appropriate configuration properties are set. The reports are saved in XML for Perfect Balance to use; they do not contain information of use to

you. They are named `${job_output_dir}/_balancer/ReduceKeyMetricList-attempt_jobid_taskid_task_attemptid.xml`. They are generated only when the counting reducer is used (that is, `oracle.hadoop.balancer.autoAnalyze=REDUCER_REPORT` when using Perfect Balance, or a call to the `Balancer.configureCountingReducer` method when using the API.

See ["Collecting Additional Metrics"](#) on page 5-7.

The reports are stored by default in the job output directory (`${mapreduce.output.fileoutputformat.outputdir}` in YARN and `${mapred.output.dir}` in MRv1). Following is the structure of that directory:

```
job_output_directory
/_SUCCESS
/_balancer
  ReduceKeyMetricList-attempt_201305031125_0016_r_000000_0.xml
  ReduceKeyMetricList-attempt_201305031125_0016_r_000001_0.xml
  .
  .
  jobanalyzer-report.html
  jobanalyzer-report.xml
  orabalancer_report.xml
/part-r-00000
/part-r-00001
  .
  .
  .
```

About Chopping

To balance a load, Perfect Balance might subpartition the values of a single reduce key and send each subpartition to a different reducer. This partitioning of values is called **chopping**.

Selecting a Chopping Method

You can configure how Perfect Balance chops the values by setting the `oracle.hadoop.balancer.enableSorting` configuration property:

- **Chopping by hash partitioning:** Set `enableSorting=false` when sorting is not required. This is the default chopping strategy.
- **Chopping by sorting:** Set `enableSorting=true` to sort the values in each subpartition and order them across all subpartitions. In any parallel sort job, each task sort the rows within the task. The job must ensure that the values in reduce task 2 are greater than values in reduce task 1, the values in reduce task 3 are greater than the values in reduce task 2, and so on. The job generates multiple files containing data in sorted order, instead of one large file with sorted data.

For example, if a key is chopped into three subpartitions, and the subpartitions are sent to reducers 5, 8 and 9, then the values for that key in reducer 9 are greater than all values for that key in reducer 8, and the values for that key in reducer 8 are greater than all values for that key in reducer 5. When `enableSorting=true`, Perfect Balance ensures this ordering across reduce tasks.

If an application requires that the data is aggregated across files, then you can disable chopping by setting `oracle.hadoop.balancer.keyLoad.minChopBytes=-1`. Perfect

Balance still offers performance gains by combining smaller reduce keys, called **bin packing**.

How Chopping Impacts Applications

If a MapReduce job aggregates the data by reduce key, then each reduce task aggregates the values for each key within that task. However, when chopping is enabled in Perfect Balance, the rows associated with a reduce key might be in different reduce tasks, leading to partial aggregation. Thus, values for a reduce key are aggregated within a reduce task, but not across reduce tasks. (The values for a reduce key across reduce tasks can be sorted, as discussed in ["Selecting a Chopping Method"](#) on page 5-13.)

When complete aggregation is required, you can disable chopping. Alternatively, you can examine the application that consumes the output of your MapReduce job. The application might work well with partial aggregation.

For example, a search engine might read in parallel the output from a MapReduce job that creates an inverted index. The output of a reduce task is a list of words, and for each word, a list of documents in which the word occurs. The word is the key, and the list of documents is the value. With partial aggregation, some words have multiple document lists instead of one aggregated list. Multiple lists are convenient for the search engine to consume in parallel. A parallel search engine might even require document lists to be split instead of aggregated into one list. See ["About the Perfect Balance Examples"](#) on page 5-17 for a Hadoop job that creates an inverted index from a document collection.

As another example, Oracle Loader for Hadoop loads data from multiple files to the correct partition of a target table. The load step is faster when there are multiple files for a reduce key, because they enable a higher degree of parallelism than loading from one file for a reduce key.

Troubleshooting Jobs Running with Perfect Balance

Check these tips for help troubleshoot a job running with Perfect Balance.

Java GC Overhead Limit Exceeded Error

If you get a Java "GC overhead limit exceeded" error on the client node while running a job with Perfect Balance, then change the client JVM garbage collector for your job.

Set the Java JVM `-XX:+UseConcMarkSweepGC` option. You can specify client JVM options before running the Hadoop job, by setting the `HADOOP_CLIENT_OPTS` variable:

```
export HADOOP_CLIENT_OPTS="-XX:+UseConcMarkSweepGC $HADOOP_CLIENT_OPTS"
```

Setting `HADOOP_CLIENT_OPTS` changes the JVM options only on the client node. It does not change JVM options in the map and reduce tasks. See the `invindx` script for an example of setting this variable.

You can also increase the client JVM heap size. See ["Java Out of Heap Space Errors"](#) on page 5-14.

Java Out of Heap Space Errors

If you get Java "out of heap space" errors on the client node while running the Perfect Balance sampler, then increase the client JVM heap size for the job.

Use the Java JVM `-Xmx` option. You can specify client JVM options before running the Hadoop job, by setting the `HADOOP_CLIENT_OPTS` variable:

```
$ export HADOOP_CLIENT_OPTS="-Xmx512M $HADOOP_CLIENT_OPTS"
```

Setting `HADOOP_CLIENT_OPTS` changes the JVM options only on the client node. It does not change JVM options in the map and reduce tasks. See the `invindx` script for an example of setting this variable.

Setting `HADOOP_CLIENT_OPTS` is sufficient to increase the heap size for the sampler, regardless of whether `oracle.hadoop.balancer.runMode` is set to `local` or `distributed`. When `runMode=local`, the sampler runs on the client node, and `HADOOP_CLIENT_OPTS` sets the heap size on the client node. When `runMode=distributed`, Perfect Balance automatically sets the heap size for the sampler Hadoop job based on the `-Xmx` setting you provide in `HADOOP_CLIENT_OPTS`. Perfect Balance never changes the heap size for the map and reduce tasks of your job, only for its sampler job.

Using the Perfect Balance API

The `oracle.hadoop.balancer.Balancer` class contains methods for creating a partitioning plan, saving the plan to a file, and running the MapReduce job using the plan. You only need to add the code to the application's job driver Java class, not redesign the application. When you run a shell script to run the application, you can include Perfect Balance configuration settings.

Modifying Your Java Code to Use Perfect Balance

The Perfect Balance installation directory contains a complete example, including input data, of a Java MapReduce program that uses the Perfect Balance API.

For a description of the inverted index example and execution instructions, see `orabalancer-2.0.0-h2/examples/invindx/README.txt`.

To explore the modified Java code, see

`orabalancer-2.0.0-h2/examples/jsrc/oracle/hadoop/balancer/examples/invindx/InvertedIndexMapred.java` or `InvertedIndexMapreduce.java`.

The modifications to run Perfect Balance include the following:

- The `createBalancer` method validates the configuration properties and returns a `Balancer` instance.
- The `waitForCompletion` method samples the data and creates a partitioning plan.
- The `addBalancingPlan` method adds the partitioning plan to the job configuration settings.
- The `configureCountingReducer` method collects additional load statistics.
- The `save` method saves the partition report and generates the Job Analyzer report.

[Example 5-4](#) shows fragments from the inverted index Java code.

Example 5-4 Running Perfect Balance in a MapReduce Job

```
.
.
.
import oracle.hadoop.balancer.Balancer;
.
.
.
```

```
///// BEGIN: CODE TO INVOKE BALANCER (PART-1, before job submission) /////
Configuration conf = job.getConfiguration();

Balancer balancer = null;

boolean useBalancer =
    conf.getBoolean("oracle.hadoop.balancer.driver.balance", true);
if(useBalancer)
{
    balancer = Balancer.createBalancer(conf);
    balancer.waitForCompletion();
    balancer.addBalancingPlan(conf);
}

if(conf.getBoolean("oracle.hadoop.balancer.tools.useCountingReducer", true))
{
    Balancer.configureCountingReducer(conf);
}
////////// END: CODE TO INVOKE BALANCER (PART-1) //////////

boolean isSuccess = job.waitForCompletion(true);

//////////
// BEGIN: CODE TO INVOKE BALANCER (PART-2, after job completion, optional)
// If balancer ran, this saves the partition file report into the _balancer
// sub-directory of the job output directory. It also writes a JobAnalyzer
// report.
Balancer.save(job);
////////// END: CODE TO INVOKE BALANCER (PART-2) //////////
.
.
.
}
```

See Also: *Oracle Big Data Appliance Perfect Balance Java API Reference*

Running Your Modified Java Code with Perfect Balance

When you run your modified Java code, you can set the Perfect Balance properties by using the standard hadoop command syntax:

```
bin/hadoop jar application_jarfile.jar ApplicationClass \
-conf application_config.xml \
-conf perfect_balance_config.xml \
-D application_config_property \
-D perfect_balance_config_property \
-libjars application_jar_path.jar...
```

[Example 5-5](#) runs a script named `pb_balanceapi.sh`, which runs the `InvertedIndexMapreduce` class example packaged in the Perfect Balance JAR file. The key load metric properties are set to the values recommended in the Job Analyzer report shown in [Figure 5-1](#).

To run the `InvertedIndexMapreduce` class example, see ["About the Perfect Balance Examples"](#) on page 5-17.

Example 5-5 *Running the InvertedIndexMapreduce Class*

```
$ cat pb_balanceapi.sh
BALANCER_HOME=/opt/oracle/orabalancer-2.0.0-h2
APP_JAR_FILE=/opt/oracle/orabalancer-2.0.0-h2/jlib/orabalancer-2.0.0.jar
```

```

export HADOOP_CLASSPATH=${BALANCER_HOME}/jlib/orabalancer-2.0.0.jar:${BALANCER_
HOME}/jlib/commons-math-2.2.jar:$HADOOP_CLASSPATH
export HADOOP_USER_CLASSPATH_FIRST=true

hadoop jar ${APP_JAR_FILE}
oracle.hadoop.balancer.examples.invindx.InvertedIndexMapreduce \
-D mapreduce.input.fileinputformat.inputdir=invindx/input \
-D mapreduce.output.fileoutputformat.outputdir=jdoe_outdir_api \
-D mapreduce.job.name=jdoe_invindx_api \
-D mapreduce.job.reduces=10 \
-D oracle.hadoop.balancer.autoAnalyze=true \
-D oracle.hadoop.balancer.linearKeyLoad.keyWeight=93.981394 \
-D oracle.hadoop.balancer.linearKeyLoad.rowWeight=0.001126 \
-D oracle.hadoop.balancer.linearKeyLoad.byteWeight=0.0

$ sh ./balanceapi.sh
14/04/14 15:03:51 INFO balancer.Balancer: Creating balancer
14/04/14 15:03:51 INFO balancer.Balancer: Starting Balancer
14/04/14 15:03:51 INFO input.FileInputFormat: Total input paths to process : 5
14/04/14 15:03:54 INFO balancer.Balancer: Balancer completed
14/04/14 15:03:55 INFO input.FileInputFormat: Total input paths to process : 5
14/04/14 15:03:55 INFO mapreduce.JobSubmitter: number of splits:5
14/04/14 15:03:55 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_
1397066986369_3510
14/04/14 15:03:55 INFO impl.YarnClientImpl: Submitted application application_
1397066986369_3510
.
.
.
File Input Format Counters
Bytes Read=112652976
File Output Format Counters
Bytes Written=384974202

```

About the Perfect Balance Examples

The Perfect Balance installation files include a full set of examples that you can run immediately. The InvertedIndex example is a MapReduce application that creates an inverted index on an input set of text files. The inverted index maps words to the location of the words in the text files. The input data is included.

About the Examples in This Chapter

The InvertedIndex example provides the basis for all examples in this chapter. They use the same data set and run the same MapReduce application. The modifications to the InvertedIndex example simply highlight the steps you must perform in running your own applications with Perfect Balance.

If you want to run the examples in this chapter, or use them as the basis for running your own jobs, then make the following changes:

- If you are modifying the examples to run your own application, then add your application JAR files to HADOOP_CLASSPATH and `-libjars`.
- Ensure that the value of `mapreduce.input.fileinputformat.inputdir` identifies the location of your data.

The `invindx/input` directory contains the sample data for the InvertedIndex example. To use this data, you must first set it up. See ["Extracting the Example Data Set"](#) on page 5-18.

- Replace `jdoe` with your Hadoop user name.
- Set the `-conf` option to an existing configuration file.

The `jdoe_conf_invindx.xml` file is a modification of the configuration file for the InvertedIndex examples. The modified file does not have performance optimizing settings. You can use the example configuration file as is or modify it. See `/opt/oracle/orabalancer-2.0.0-h2/examples/invindx/conf_mapreduce.xml` (or `conf_mapred.xml`).
- Review the configuration settings in the file and in the shell script to ensure they are appropriate for your job.
- You can run the browser from your laptop or connect to Oracle Big Data Appliance using a client that supports graphical interfaces, such as VNC.

Extracting the Example Data Set

To run the InvertedIndex examples or any of the examples in this chapter, you must first set up the data files.

To extract the InvertedIndex data files:

1. Log in to a server where Perfect Balance is installed.
2. Change to the `examples/invindx` subdirectory:

```
cd /opt/oracle/orabalancer-2.0.0-h2/examples/invindx
```
3. Unzip the data and copy it to the HDFS `invindx/input` directory:

```
./invindx -setup
```

For complete instructions for running the InvertedIndex example, see `/opt/oracle/orabalancer-2.0.0-h2/examples/invindx/README.txt`.

Perfect Balance Configuration Property Reference

This section describes the Perfect Balance configuration properties and a few generic Hadoop MapReduce properties that Perfect Balance reads from the job configuration:

- [MapReduce Configuration Properties](#)
- [Job Analyzer Configuration Properties](#)
- [Perfect Balance Configuration Properties](#)

See "[About Configuring Perfect Balance](#)" on page 5-9 for a list of the properties organized into functional categories.

Note: CDH5 deprecates many MapReduce properties and replaces them with new properties. Perfect Balance continues to work with the old property names, but Oracle recommends that you use the new names. For the new MapReduce property names, see the Cloudera website at:

<http://archive.cloudera.com/cdh5/cdh/5/hadoop/hadoop-project-dist/hadoop-common/DeprecatedProperties.html>

MapReduce Configuration Properties

mapreduce.input.fileinputformat.inputdir

Type: String

Default Value: Not defined

Description: A comma-separated list of input directories.

mapreduce.inputformat.class

Type: String

Default Value: `org.apache.hadoop.mapreduce.lib.input.TextInputFormat`

Description: The full name of the `InputFormat` class.

mapreduce.map.class

Type: String

Default Value: `org.apache.hadoop.mapreduce.Mapper`

Description: The full name of the mapper class.

mapreduce.output.fileoutputformat.outputdir

Type: String

Default Value: Not defined

Description: The job output directory.

mapreduce.partition.class

Type: String

Default Value: `org.apache.hadoop.mapreduce.lib.partition.HashPartitioner`

Description: The full name of the partitioner class.

mapreduce.reduce.class

Type: String

Default Value: `org.apache.hadoop.mapreduce.Reducer`

Description: The full name of the reducer class.

Job Analyzer Configuration Properties

oracle.hadoop.balancer.application_id

Type: String

Default Value: Not defined

Description: The job identifier of the job you want to analyze with Job Analyzer. This property is a parameter to the Job Analyzer utility in standalone mode on YARN clusters; it does not apply to MRv1 clusters. See ["Running Job Analyzer as a Standalone Utility"](#) on page 5-4.

oracle.hadoop.balancer.tools.jobConfPath

Type: String

Default Value: `${mapreduce.output.fileoutputformat.outputdir}/_logs/history`

Description: The path to a Hadoop job configuration file in MRv1 (not applicable to YARN). Job Analyzer uses this setting to locate the file.

oracle.hadoop.balancer.tools.jobHistoryPath**Type:** String**Default Value:** `${mapreduce.output.fileoutputformat.outputdir}/_logs/history`**Description:** The path to a Hadoop job history file in MRv1 (not applicable to YARN). Job Analyzer uses this setting to locate the file.**oracle.hadoop.balancer.tools.writeKeyBytes****Type:** Boolean**Default Value:** `false`**Description:** Controls whether the counting reducer collects the byte representations of the reduce keys for the Job Analyzer. Set this property to `true` to represent the unique key values in Base64 encoding in the report. A string representation of the key, created using `key.toString`, is also provided in the report. This string value may not be unique for each key.**Perfect Balance Configuration Properties****oracle.hadoop.balancer.autoAnalyze****Type:** Enumeration**Default Value:** `BASIC_REPORT` if `oracle.hadoop.balancer.autoBalance` is `true`; otherwise `NONE`**Description:** Controls the behavior of the Job Analyzer when it is called using Perfect Balance. The following values are valid:

- `NONE`: Disables Job Analyzer.
- `BASIC_REPORT`: Enables Job Analyzer
- `REDUCER_REPORT`: Enables Job Analyzer such that it collects additional load statistics for each reduce task in a job. See ["Collecting Additional Metrics"](#) on page 5-7..

Perfect Balance uses this property; the Perfect BalanceAPI does not.

oracle.hadoop.balancer.autoBalance**Type:** Boolean**Default Value:** `false`**Description:** Controls whether load balancing is enabled. Set to `false` to turn off balancing. Perfect Balance uses this property; the Perfect Balance API does not.**oracle.hadoop.balancer.confidence****Type:** Float**Default Value:** `0.95`**Description:** The statistical confidence indicator for the load factor specified by the `oracle.hadoop.balancer.maxLoadFactor` property.This property accepts values greater than or equal to 0.5 and less than 1.0 ($0.5 \leq \text{value} < 1.0$). A value less than 0.5 resets the property to its default value. Oracle recommends a value greater than or equal to 0.9. Typical values are 0.95 and 0.99.**oracle.hadoop.balancer.enableSorting****Type:** Boolean**Default Value:** `false`

Description: Controls how the map output keys are **chopped**, that is, split into smaller keys:

- `false`: Uses a hash function.
- `true`: Uses the map output key sorting comparator as a total-order partitioning function. The balancer preserves the total order over the values of the chopped keys.

Perfect Balance sets this property to `false` if the job is not configured for secondary sorting.

oracle.hadoop.balancer.inputFormat.mapred.map.tasks

Type: Integer

Default Value: 100

Description: Sets the Hadoop `mapred.map.tasks` property for the duration of sampling, just before calling the input format `getSplits` method. It does not change `mapred.map.tasks` for the actual job. The optimal number of map tasks is a trade-off between obtaining a good sample (larger number) and having finite memory resources (smaller number).

Set this property to a value greater than or equal to one (1). A value less than 1 disables the property.

Some input formats, such as `DBInputFormat`, use this property as a hint to determine the number of splits returned by `getSplits`. Higher values indicate that more chunks of data are sampled at random, which improves the sample.

You can increase the value for larger data sets, that is, more than a million rows of about 100 bytes per row. However, extremely large values can cause the input format's `getSplits` method to run out of memory by returning too many splits.

oracle.hadoop.balancer.inputFormat.mapred.max.split.size

Type: Long

Default Value: 1048576 (1 MB)

Description: Sets the Hadoop `mapred.max.split.size` property for the duration of sampling, just before calling the input format's `getSplits` method. It does not change `mapred.max.split.size` for the actual job.

Set this property to a value greater than or equal to one (1). A value less than 1 disables the property. The optimal split size is a trade-off between obtaining a good sample (smaller splits) and efficient I/O performance (larger splits).

Some input formats, such as `FileInputFormat`, use the maximum split size as a hint to determine the number of splits returned by `getSplits`. Smaller split sizes indicate that more chunks of data are sampled at random, which improves the sample. Set the value small enough for good sampling performance, but no smaller. Extremely small values can cause inefficient I/O performance, while not improving the sample.

You can increase the value for larger data sets (tens of terabytes) or if the input format's `getSplits` method throws an out of memory error. Large splits are better for I/O performance, but not for sampling.

oracle.hadoop.balancer.keyLoad.minChopBytes

Type: Long

Default Value: 0

Description: Controls whether Perfect Balance chops large map output keys into medium keys:

- -1: Perfect Balance does not chop large map output keys.
- 0: Perfect Balance chops large map output keys and determines the optimal size of each medium key.
- *Positive integer*: Perfect Balance chops large map output keys into medium keys with a size greater than or equal to the specified integer.

oracle.hadoop.balancer.linearKeyLoad.byteWeight

Type: Float

Default Value: 0.05

Description: Weights the number of bytes per key in the linear key load model specified by the `oracle.hadoop.balancer.KeyLoadLinear` class.

oracle.hadoop.balancer.linearKeyLoad.feedbackDir

Type: String

Default Value: Not defined

Description: The path to a directory that contains the Job Analyzer report for a job that it previously analyzed. The sampler reads this report for feedback to use to optimize the current balancing plan. You can set this property to the Job Analyzer report directory of a job that is the same or similar to the current job, so that the feedback is directly applicable.

If the feedback directory contains a Job Analyzer report with recommended values for the Perfect Balance linear key load model coefficients, then Perfect Balance automatically reads and uses them. The recommended values take precedence over user-specified values in these configuration parameters:

- `oracle.hadoop.balancer.linearKeyLoad.byteWeight`
- `oracle.hadoop.balancer.linearKeyLoad.keyWeight`
- `oracle.hadoop.balancer.linearKeyLoad.rowWeight`

Job Analyzer attempts to recommend good values for these coefficients. However, Perfect Balance reads the load model coefficients from this list of configuration properties under the following circumstances:

- The `feedbackDir` property is not set.
- The `feedbackDir` property is set, but the Job Analyzer report in the specified directory does not contain a good recommendation for the load model coefficients.

oracle.hadoop.balancer.linearKeyLoad.keyWeight

Type: Float

Default Value: 50.0

Description: Weights the number of medium keys per large key in the linear key load model specified by the `oracle.hadoop.balancer.KeyLoadLinear` class.

oracle.hadoop.balancer.linearKeyLoad.rowWeight

Type: Float

Default Value: 0.05

Description: Weights the number of rows per key in the linear key load model specified by the `oracle.hadoop.balancer.KeyLoadLinear` class.

oracle.hadoop.balancer.maxLoadFactor**Type:** Float**Default Value:** 0.05

Description: The target reducer load factor that you want the balancer's partition plan to achieve.

The load factor is the relative deviation from an estimated value. For example, if `maxLoadFactor=0.05` and `confidence=0.95`, then with a confidence greater than 95%, the job's reducer loads should be, at most, 5% greater than the value in the partition plan.

The values of these two properties determine the sampler's stopping condition. The balancer samples until it can generate a plan that guarantees the specified load factor at the specified confidence level. This guarantee may not hold if the sampler stops early because of other stopping conditions, such as the number of samples exceeds [oracle.hadoop.balancer.maxSamplesPct](#). The partition report logs the stopping condition.

See [oracle.hadoop.balancer.confidence](#).

oracle.hadoop.balancer.maxSamplesPct**Type:** Float**Default Value:** 0.01 (1%)

Description: Limits the number of samples that Perfect Balance can collect to a fraction of the total input records. A value less than zero disables the property (no limit).

You may need to increase the value for Hadoop applications with very unbalanced reducer partitions or densely clustered map-output keys. The sampler needs to sample more data to achieve a good partitioning plan in these cases.

See [oracle.hadoop.balancer.useClusterStats](#).

oracle.hadoop.balancer.minSplits**Type:** Integer**Default Value:** 5

Description: Sets the minimum number of splits that the sampler reads. If the total number of splits is less than this value, then the sampler reads all splits. Set this property to a value greater than or equal to one (1). A nonpositive number sets the property to 1.

oracle.hadoop.balancer.numThreads**Type:** Integer**Default Value:** 5

Description: Number of sampler threads. Set this value based on the processor and memory resources available on the node where the job is initiated. A higher number of sampler threads implies higher concurrency in sampling. Set this property to one (1) to disable multithreading in the sampler.

oracle.hadoop.balancer.report.overwrite**Type:** Boolean**Default Value:** false

Description: Controls whether Perfect Balance overwrites files in the location specified by the [oracle.hadoop.balancer.reportPath](#) property. By default, Perfect Balance

does not overwrite files; it throws an exception. Set this property to `true` to allow partition reports to be overwritten.

oracle.hadoop.balancer.reportPath

Type: String

Default Value: `directory/orabalancer_report-random_unique_string.xml`, where `directory` for HDFS is the home directory of the user who submits the job. For the local file system, it is the directory where the job is submitted.

Description: The path where Perfect Balance writes the partition report before the Hadoop job output directory is available, that is, before the MapReduce job finishes running. At the end of the job, Perfect Balance moves the file to `job_output_dir/_balancer/orabalancer_report.xml`. In the API, the `save` method does this task.

oracle.hadoop.balancer.runMode

Type: String

Default Value: `local`

Description: Specifies how to run the Perfect Balance sampler. The following values are valid:

- `local`: The sampler runs on the client node where the job is submitted.
- `distributed`: The sampler runs as a Hadoop job. If the job uses the distributed cache, then Perfect Balance automatically sets this property to `distributed`.

If this property is set to an invalid string, Perfect Balance resets it to `local`.

oracle.hadoop.balancer.tmpDir

Type: String

Default Value: `/tmp/orabalancer-user_name`

Description: The path to a staging directory in the file system of the job output directory (HDFS or local). Perfect Balance creates the directory if it does not exist, and copies the partition report to it for loading into the Hadoop distributed cache.

oracle.hadoop.balancer.useClusterStats

Type: Boolean

Default Value: `true`

Description: Enables the sampler to use cluster sampling statistics. These statistics improve the accuracy of sampled estimates, such as the number of records in a map-output key, when the map-output keys are distributed in clusters across input splits, instead of being distributed independently across all input splits.

Set this property to `false` only if you are absolutely certain that the map-output keys are not clustered. This setting improves the sampler's estimates only when there is, in fact, no clustering. Oracle recommends leaving this property set to `true`, because the distribution of map-output keys is usually unknown.

oracle.hadoop.balancer.useMapreduceApi

Type: Boolean

Default Value: `true`

Description: Identifies the MapReduce API used in the Hadoop job:

- `true`: The job uses the `mapreduce` API.
- `false`: The job uses the `mapred` API.

Glossary

Apache Flume

A distributed service for collecting and aggregating data from almost any source into a data store such as HDFS or HBase.

See also [Apache HBase](#); [HDFS](#).

Apache HBase

An open-source, column-oriented database that provides random, read/write access to large amounts of sparse data stored in a CDH cluster. It provides fast lookup of values by key and can perform thousands of insert, update, and delete operations per second.

Apache Hive

An open-source data warehouse in CDH that supports data summarization, ad hoc querying, and data analysis of data stored in HDFS. It uses a SQL-like language called HiveQL. An interpreter generates MapReduce code from the HiveQL queries.

By using Hive, you can avoid writing MapReduce programs in Java.

See also [Hive Thrift](#); [HiveQL](#); [MapReduce](#).

Apache Sentry

Integrates with the Hive and Impala SQL-query engines to provide fine-grained authorization to data and metadata stored in Hadoop.

Apache Solr

Provides an enterprise search platform that includes full-text search, faceted search, geospatial search, and hit highlighting.

Apache Spark

A fast engine for processing large-scale data. It supports Java, Scala, and Python applications. Because it provides primitives for in-memory cluster computing, it is particularly suited to machine-learning algorithms. It promises performance up to 100 times faster than MapReduce.

Apache Sqoop

A command-line tool that imports and exports data between HDFS or Hive and structured databases. The name Sqoop comes from "SQL to Hadoop." Oracle R Advanced Analytics for Hadoop uses the Sqoop executable to move data between HDFS and Oracle Database.

Apache YARN

An updated version of MapReduce, also called MapReduce 2. The acronym stands for Yet Another Resource Negotiator.

ASR

Oracle Auto Service Request, a software tool that monitors the health of the hardware and automatically generates a service request if it detects a problem.

See also [OASM](#).

Balancer

A service that ensures that all nodes in the cluster store about the same amount of data, within a set range. Data is balanced over the nodes in the cluster, not over the disks in a node.

CDH

Cloudera's Distribution including Apache Hadoop, the version of Apache Hadoop and related components installed on Oracle Big Data Appliance.

Cloudera Hue

Hadoop User Experience, a web user interface in CDH that includes several applications, including a file browser for HDFS, a job browser, an account management tool, a MapReduce job designer, and Hive wizards. Cloudera Manager runs on Hue.

See also [HDFS](#); [Apache Hive](#).

Cloudera Impala

A massively parallel processing query engine that delivers better performance for SQL queries against data in HDFS and HBase, without moving or transforming the data.

Cloudera Manager

Cloudera Manager enables you to monitor, diagnose, and manage CDH services in a cluster.

The Cloudera Manager agents on Oracle Big Data Appliance also provide information to Oracle Enterprise Manager, which you can use to monitor both software and hardware.

Cloudera Navigator

Verifies access privileges and audits access to data stored in Hadoop, including Hive metadata and HDFS data accessed through HDFS, Hive, or HBase.

Cloudera Search

Provides search and navigation tools for data stored in Hadoop.

Cloudera's Distribution including Apache Hadoop (CDH)

See [CDH](#).

cluster

A group of servers on a network that are configured to work together. A server is either a master node or a worker node.

All servers in an Oracle Big Data Appliance rack form a cluster. Servers 1, 2, and 3 are master nodes. Servers 4 to 18 are worker nodes.

See [Hadoop](#).

DataNode

A server in a CDH cluster that stores data in HDFS. A DataNode performs file system operations assigned by the NameNode.

See also [HDFS](#); [NameNode](#).

Flume

See [Apache Flume](#).

Hadoop

A batch processing infrastructure that stores files and distributes work across a group of servers. Oracle Big Data Appliance uses Cloudera's Distribution including Apache Hadoop (CDH).

Hadoop Distributed File System (HDFS)

See [HDFS](#).

Hadoop User Experience (Hue)

See [Cloudera Hue](#).

HBase

See [Apache HBase](#).

HDFS

Hadoop Distributed File System, an open-source file system designed to store extremely large data files (megabytes to petabytes) with streaming data access patterns. HDFS splits these files into data blocks and distributes the blocks across a CDH cluster.

When a data set is larger than the storage capacity of a single computer, then it must be partitioned across several computers. A distributed file system can manage the storage of a data set across a network of computers.

See also [cluster](#).

Hive

See [Apache Hive](#).

Hive Thrift

A remote procedure call (RPC) interface for remote access to CDH for Hive queries.

See also [CDH](#); [Apache Hive](#).

HiveQL

A SQL-like query language used by Hive.

See also [Apache Hive](#).

HotSpot

A Java Virtual Machine (JVM) that is maintained and distributed by Oracle. It automatically optimizes code that executes frequently, leading to high performance. HotSpot is the standard JVM for the other components of the Oracle Big Data Appliance stack.

Hue

See [Cloudera Hue](#).

Impala

See [Cloudera Impala](#).

Java HotSpot Virtual Machine

See [HotSpot](#).

JobTracker

A service that assigns tasks to specific nodes in the CDH cluster, preferably those nodes storing the data. MRv1 only.

See also [Hadoop](#); [MapReduce](#).

Kerberos

A network authentication protocol that helps prevent malicious impersonation. It was developed at the Massachusetts Institute of Technology (MIT).

Mahout

Apache Mahout is a machine learning library that includes core algorithms for clustering, classification, and batch-based collaborative filtering.

MapReduce

A parallel programming model for processing data on a distributed system. Two versions of MapReduce are available, MapReduce 1 and YARN (MapReduce 2). The default version on Oracle Big Data Appliance 3.0 and later is YARN.

A MapReduce program contains these functions:

- Mappers: Process the records of the data set.
- Reducers: Merge the output from several mappers.
- Combiners: Optimizes the result sets from the mappers before sending them to the reducers (optional and not supported by all applications).

See also [Apache YARN](#).

MySQL Database

A SQL-based relational database management system. Cloudera Manager, Oracle Data Integrator, Hive, and Oozie use MySQL Database as a metadata repository on Oracle Big Data Appliance.

NameNode

A service that maintains a directory of all files in HDFS and tracks where data is stored in the CDH cluster.

See also [HDFS](#).

Navigator

See [Cloudera Navigator](#).

node

A server in a CDH cluster.

See also [cluster](#).

NodeManager

A service that runs on each node and executes the tasks assigned to it by the ResourceManager. YARN only.

See also [ResourceManager](#); [YARN](#).

NoSQL Database

See [Oracle NoSQL Database](#).

OASM

Oracle Automated Service Manager, a service for monitoring the health of Oracle Sun hardware systems. Formerly named Sun Automatic Service Manager (SASM).

Oozie

An open-source workflow and coordination service for managing data processing jobs in CDH.

Oracle Database Instant Client

A small-footprint client that enables Oracle applications to run without a standard Oracle Database client.

Oracle Linux

An open-source operating system. Oracle Linux 5.6 is the same version used by Exalogic 1.1. It features the Oracle Unbreakable Enterprise Kernel.

Oracle NoSQL Database

A distributed key-value database that supports fast querying of the data, typically by key lookup.

Oracle R Distribution

An Oracle-supported distribution of the R open-source language and environment for statistical analysis and graphing.

Oracle R Enterprise

A component of the Oracle Advanced Analytics Option. It enables R users to run R commands and scripts for statistical and graphical analyses on data stored in an Oracle database.

Pig

An open-source platform for analyzing large data sets that consists of the following:

- Pig Latin scripting language
- Pig interpreter that converts Pig Latin scripts into MapReduce jobs

Pig runs as a client application.

See also [MapReduce](#).

Puppet

A configuration management tool for deploying and configuring software components across a cluster. The Oracle Big Data Appliance initial software installation uses Puppet.

The Puppet tool consists of these components: puppet agents, typically just called puppets; the puppet master server; a console; and a cloud provisioner.

See also [puppet agent](#); [puppet master](#).

puppet agent

A service that primarily pulls configurations from the puppet master and applies them. Puppet agents run on every server in Oracle Big Data Appliance.

See also [Puppet](#); [puppet master](#)

puppet master

A service that primarily serves configurations to the puppet agents.

See also [Puppet](#); [puppet agent](#).

ResourceManager

A service that assigns tasks to specific nodes in the CDH cluster, preferably those nodes storing the data. YARN only.

See also [Hadoop](#); [YARN](#).

Search

See [Cloudera Search](#).

Sentry

See [Apache Sentry](#).

Solr

See [Apache Solr](#).

Spark

See [Apache Spark](#).

Sqoop

See [Apache Sqoop](#).

table

In Hive, all files in a directory stored in HDFS.

See also [HDFS](#).

TaskTracker

A service that runs on each node and executes the tasks assigned to it by the JobTracker service. MRv1 only.

See also [JobTracker](#).

Whirr

Apache Whirr is a set of libraries for running cloud services.

YARN

See [Apache YARN](#).

ZooKeeper

A MapReduce 1 centralized coordination service for CDH distributed processes that maintains configuration information and naming, and provides distributed synchronization and group services.

Index

A

- activity reports, 2-28
- Apache Sentry, 2-25
- application adapters, 1-8
- applications
 - data pull, 4-1
 - data push, 4-2
- Audit Vault plug-in, 2-28
- auditing data collected from services, 2-27
- authentication, 3-1
- authorization, 2-25
- autoAnalyze configuration property, 5-11, 5-20
- autoAnalyze property, 5-7
- autoBalance configuration property, 5-10, 5-20
- Automated Service Manager
 - See* OASM

B

- BALANCER_HOME environment variable, 5-3
- bdadiag utility, 2-29
- Berkeley DB, 1-5
- best practices, 5-1
- big data description, 1-1
- business intelligence, 1-3, 1-5, 1-9
- byteWeight configuration property, 5-22

C

- CDH
 - about, 1-3
 - diagnostics, 2-29
 - file system, 1-5
 - remote client access, 3-2
 - security, 3-1
 - version, 2-9
- chopped keys, 5-21
- chunking files, 1-5
- client access
 - HDFS cluster, 3-4
 - HDFS secured cluster, 3-5
 - Hive, 3-8
- client configuration, 3-2
- Cloudera Manager
 - about, 2-3

- accessing administrative tools, 2-4
- connecting to, 2-3
- effect of hardware failure on, 2-17
- software dependencies, 2-17
- starting, 2-3
- UI overview, 2-3
- version, 2-9
- Cloudera's Distribution including Apache Hadoop
 - See* CDH
- clusters, definition, 1-3
- confidence configuration property, 5-20
- Counting Reducer, 5-2
- CSV files, 1-7

D

- Data Pump files, 1-8
- data replication, 1-5
- data skew, 5-1
- DataNode, 2-16
- dba group, 2-25
- diagnostics, collecting, 2-29
- disks, 2-15
- dnsmasq service, 4-4
- duplicating data, 1-5

E

- emcli utility, 2-2
- enableSorting configuration property, 5-20
- encryption, 2-25
- engineered systems, 1-3
- Exadata Database Machine, 1-3
- Exadata InfiniBand connections, 4-2
- Exalytics In-Memory Machine, 1-3
- external tables, 1-7

F

- failover
 - JobTracker, 2-14
 - NameNode, 2-13
- feedbackDir configuration property, 5-22
- files, recovering HDFS, 3-11
- first NameNode, 2-16
- Flume, 2-10, 2-24

ftp.oracle.com, 2-29

G

GC overhead limit exceeded error, 5-14
groups, 2-24, 3-9

H

Hadoop Distributed File System
 See HDFS
hadoop group, 3-9
Hadoop version, 1-3
HADOOP_CLASSPATH environment variable, 5-3, 5-17
HADOOP_USER_CLASSPATH_FIRST environment variable, 5-3
HBase, 2-10, 2-24
HDFS
 about, 1-3, 1-5
 auditing, 2-27
 data files, 1-7
 user identity, 2-24
help from Oracle Support, 2-29
Hive, 2-24
 about, 1-5
 auditing, 2-27
 client access, 3-8
 node location, 2-17
 software dependencies, 2-17
 tables, 3-9
 user identity, 2-24
hive group, 3-9
HiveQL, 1-5
HotSpot
 See Java HotSpot Virtual Machine
Hue, 2-17
 user identity, 2-24
 users, 3-9

I

Impala, 2-10
InfiniBand connections to Exadata, 4-2
InfiniBand network configuration, 4-1
inputFormat.mapred.* configuration properties, 5-21
installing CDH client, 3-2

J

Java HotSpot Virtual Machine, 2-9
JDBC client, configuring for SDP, 4-5
Job Analyzer, 5-2, 5-4
job duration, 5-1
jobconfPath property, 5-19
jobHistoryPath configuration property, 5-20
JobTracker
 failover, 2-14
 security, 3-1
 user identity, 2-24

Index-2

JobTracker node, 2-17

K

Kerberos authentication, 3-1
Kerberos commands, 3-1
Kerberos user setup, 3-10
key chopping, 5-2
keyLoad.minChopBytes configuration property, 5-21
keys, assigning to reducers, 5-2
key-value database, 1-5
keyWeight configuration property, 5-22
knowledge modules, 1-8

L

linearKeyLoad properties, 5-8
linearKeyLoad.* configuration properties, 5-22
Linux
 disk location, 2-16
 installation, 2-8
load, 5-1
Load Balancer, 5-2
loading data, 1-7, 1-8
login privileges, 3-11

M

mapper workload, 5-2
mapred configuration properties, 5-19
mapred user, 2-24
mapred.map.tasks configuration property, 5-21
MapReduce, 1-4, 1-6, 2-27, 3-1, 3-9
mapreduce configuration properties, 5-19
map.tasks property, 5-21
maxLoadFactor configuration property, 5-23
maxSamplesPct configuration property, 5-23
max.split.size configuration property, 5-21
minChopBytes configuration property, 5-21
minSplits configuration property, 5-23
monitoring activity, 2-28
multirack clusters
 service locations, 2-11
MySQL Database
 about, 2-17
 port number, 2-26
 user identity, 2-24
 version, 2-9

N

NameNode, 3-1
 first, 2-16
NameNode failover, 2-13
Navigator, 2-10
NoSQL databases
 See Oracle NoSQL Database
numThreads configuration property, 5-23

O

OASM, port number, 2-26
ODI
 See Oracle Data Integrator
oinstall group, 2-25, 3-9
on-disk encryption, 2-25
Oozie, 2-18
 auditing, 2-27
 software dependencies, 2-17
 software services, 2-25
 user identity, 2-25
openib.conf file, 4-4
operating system users, 2-24
Oracle Audit Vault and Database Firewall, 2-27
Oracle Automated Service Manager
 See OASM
Oracle Data Integrator
 about, 1-8
 node location, 2-17
 software dependencies, 2-17
 version, 2-9
Oracle Data Integrator agent, 2-26
Oracle Data Pump files, 1-8
Oracle Database Instant Client, 2-9
Oracle Exadata Database Machine, 1-3, 4-1
 using as a CDH client, 3-3
Oracle Exalytics In-Memory Machine, 1-3
Oracle Linux
 about, 1-3
 relationship to HDFS, 1-4
 version, 2-9
Oracle Loader for Hadoop, 1-8, 2-9
Oracle NoSQL Database
 about, 1-5, 1-8
 port numbers, 2-26
 version, 2-9
Oracle R Advanced Analytics for Hadoop, 1-8, 2-9
Oracle R Enterprise, 1-9
Oracle SQL Connector for HDFS, 1-7
Oracle Support, creating a service request, 2-29
oracle user, 2-25, 3-9
Oracle XQuery for Hadoop, 1-8, 2-9
oracle.hadoop.balancer.* configuration
 properties, 5-20
oracle.hadoop.balancer.autoAnalyze configuration
 property, 5-11
oracle.hadoop.balancer.autoAnalyze property, 5-7
oracle.hadoop.balancer.autoBalance configuration
 property, 5-10
oracle.hadoop.balancer.Balancer class, 5-15
oracle.hadoop.balancer.KeyLoadLinear class, 5-22
oracle.hadoop.balancer.linearKeyLoad.*
 properties, 5-8
out of heap space errors, 5-14

P

partitioning, 2-15, 5-2
passphrase-based encryption, 2-25
Perfect Balance

 application requirements, 5-3
 basic steps, 5-3
 description, 5-1
planning applications, 1-3
port map, 2-26
port numbers, 2-26
pulling data into Exadata, 4-1
puppet
 port numbers, 2-26
 security, 2-26
 user identity, 2-25
puppet master
 node location, 2-17
pushing data into Exadata, 4-2

R

R Connector
 See Oracle R Advanced Analytics for Hadoop
R distribution, 2-9
R language support, 1-8
range partitioning, 5-2
recovering HDFS files, 3-11
reducer load, 5-1
remote client access, 3-2, 3-8
replicating data, 1-5
report.overwrite configuration property, 5-23
reportPath configuration property, 5-24
rowWeight configuration property, 5-22
rpc.statd service, 2-26

S

SDP listener configuration, 4-5
SDP over InfiniBand, 4-1
SDP, enabling on Exadata, 4-4
Search, 2-10
security, 2-24
Sentry, 2-25
service requests, creating for CDH, 2-29
service tags, 2-26
services
 auditing, 2-27
 node locations, 2-11
skew, 5-1
Sockets Direct Protocol, 4-1
software components, 2-8
software framework, 1-3
software services
 node locations, 2-11
 port numbers, 2-26
Spark, 2-10
Sqoop, 2-10, 2-25
ssh service, 2-26
svctag user, 2-25

T

tables, 1-7, 1-8, 3-9
TaskTracker
 user identity, 2-24

- tmpDir configuration property, 5-24
- tools.* configuration properties, 5-19
- TPM encryption, 2-25
- trash facility, 3-11
- trash facility, disabling, 3-13
- trash interval, 3-12
- troubleshooting CDH, 2-29

U

- uploading diagnostics, 2-29
- useClusterStats configuration property, 5-24
- useMapreduceApi configuration property, 5-24
- user accounts, 3-9
- user groups, 3-9
- users
 - Cloudera Manager, 2-4
 - operating system, 2-24

W

- writeKeyBytes configuration property, 5-20

X

- xinetd service, 2-26
- XQuery connector
 - See* Oracle XQuery for Hadoop

Y

- YARN support, 1-7, 2-22

Z

- ZooKeeper, 2-25